

云容器实例

API 参考

文档版本 01
发布日期 2023-03-30



版权所有 © 华为云计算技术有限公司 2024。保留一切权利。

非经本公司书面许可，任何单位和个人不得擅自摘抄、复制本文档内容的部分或全部，并不得以任何形式传播。

商标声明



HUAWEI和其他华为商标均为华为技术有限公司的商标。

本文档提及的其他所有商标或注册商标，由各自的所有人拥有。

注意

您购买的产品、服务或特性等应受华为云计算技术有限公司商业合同和条款的约束，本文档中描述的全部或部分产品、服务或特性可能不在您的购买或使用范围之内。除非合同另有约定，华为云计算技术有限公司对本文档内容不做任何明示或暗示的声明或保证。

由于产品版本升级或其他原因，本文档内容会不定期进行更新。除非另有约定，本文档仅作为使用指导，本文档中的所有陈述、信息和建议不构成任何明示或暗示的担保。

目录

1 使用前必读	1
1.1 概述	1
1.2 调用说明	1
1.3 终端节点	1
1.4 约束与限制	1
1.5 基本概念	1
1.6 API 版本选择建议	2
2 如何调用 API	4
2.1 构造请求	4
2.2 认证鉴权	7
2.3 返回结果	9
3 快速入门	11
4 API	15
4.1 Metrics	15
4.1.1 获取指定 namespace 下所有 pods 的 metrics 信息	15
4.1.2 获取指定 namespace 下指定 pod 的 metrics 信息	32
4.2 EIPPool	45
4.2.1 查询指定 namespace 下的 EIP Pools	45
4.2.2 创建 EIPPool	62
4.2.3 删除指定的 EIPPool	104
4.2.4 查询指定的 EIPPool	117
4.2.5 更新 EIPPool	129
4.2.6 替换 EIPPool	142
4.2.7 查询指定的 EIPPool	177
4.3 Network	189
4.3.1 删除指定 namespace 下的 Networks	189
4.3.2 查询指定 namespace 下的 Networks	201
4.3.3 创建 Network	216
4.3.4 删除 Network	251
4.3.5 查询 Network	272
4.3.6 查询 Network 状态	281
5 Kubernetes API	292

5.1 ConfigMap.....	292
5.1.1 删除指定 namespace 下的 ConfigMaps.....	292
5.1.2 查询指定 namespace 下的 ConfigMaps.....	304
5.1.3 创建 ConfigMap.....	319
5.1.4 删除 ConfigMap.....	352
5.1.5 查询 ConfigMap.....	363
5.1.6 更新 ConfigMap.....	372
5.1.7 替换 ConfigMap.....	383
5.2 Pod.....	410
5.2.1 删除指定 namespace 下的 Pods.....	410
5.2.2 查询指定 namespace 下的 Pods.....	422
5.2.3 创建 Pod.....	502
5.2.4 删除 Pod.....	825
5.2.5 查询 Pod.....	973
5.2.6 更新 Pod.....	1048
5.2.7 替换 Pod.....	1124
5.2.8 进入容器执行命令.....	1368
5.2.9 进入容器执行命令.....	1370
5.2.10 查询 Pod 日志.....	1373
5.2.11 查询 Pod 状态.....	1376
5.2.12 查询用户所有的 Pods.....	1450
5.3 StorageClass.....	1529
5.3.1 查询/apis/storage.k8s.io/v1 版本的所有 storage class.....	1529
5.3.2 查询指定的 storage class.....	1544
5.4 Service.....	1553
5.4.1 查询指定 namespace 下的 Services.....	1553
5.4.2 创建 Service.....	1574
5.4.3 删除 Service.....	1635
5.4.4 查询 Service.....	1647
5.4.5 更新 Service.....	1662
5.4.6 替换 Service.....	1679
5.4.7 查询 service 状态.....	1727
5.5 Deployment.....	1742
5.5.1 查询用户所有 Deployments.....	1742
5.5.2 删除指定 namespace 下 Deployments.....	1821
5.5.3 查询指定 namespace 下的 Deployments.....	1833
5.5.4 创建 Deployment.....	1912
5.5.5 删除 Deployment.....	2227
5.5.6 查询 Deployment.....	2240
5.5.7 更新 Deployment.....	2313
5.5.8 替换 Deployment.....	2388
5.5.9 查询 Deployment 的伸缩操作.....	2629

5.5.10 更新 Deployment 的伸缩操作.....	2638
5.5.11 替换 Deployment 的伸缩操作.....	2649
5.5.12 查询 Deployment 状态.....	2675
5.6 Ingress.....	2748
5.6.1 删除指定 namespace 下的 ingresses.....	2748
5.6.2 查询指定 namespace 下的 Ingresses.....	2761
5.6.3 创建 Ingress.....	2780
5.6.4 删除 Ingress.....	2832
5.6.5 查询 Ingress.....	2844
5.6.6 更新 Ingress.....	2857
5.6.7 替换 Ingress.....	2874
5.6.8 查询 Ingress 状态.....	2915
5.7 OpenAPIv2.....	2928
5.7.1 查询 open api swagger 信息.....	2928
5.8 VolcanoJob.....	2930
5.8.1 删除指定 namespace 下的 Volcano Jobs.....	2930
5.8.2 查询指定 namespace 下的 Volcano Jobs.....	2941
5.8.3 创建 Volcano Job.....	2962
5.8.4 删除 Volcano Job.....	3023
5.8.5 查询 Volcano Job 详情.....	3035
5.8.6 更新 Volcano Job.....	3052
5.8.7 替换 Volcano Job.....	3070
5.9 Namespace.....	3118
5.9.1 查询所有 Namespaces.....	3118
5.9.2 创建 Namespace.....	3133
5.9.3 删除 Namespace.....	3168
5.9.4 查询 Namespace.....	3179
5.10 ClusterRole.....	3188
5.10.1 获取 ClusterRole 列表.....	3188
5.10.2 查询指定的 ClusterRole.....	3203
5.11 Secret.....	3214
5.11.1 删除指定 namespace 下的 Secrets.....	3214
5.11.2 查询指定 namespace 下的 Secrets.....	3227
5.11.3 创建 Secret.....	3240
5.11.4 删除 Secret.....	3274
5.11.5 查询 Secret.....	3286
5.11.6 更新 Secret.....	3295
5.11.7 替换 Secret.....	3306
5.12 Endpoint.....	3333
5.12.1 查询指定 namespace 下的 Endpoints.....	3333
5.12.2 创建 Endpoint.....	3348
5.12.3 删除 Endpoint.....	3390

5.12.4 查询 Endpoint.....	3402
5.12.5 更新 Endpoint.....	3414
5.12.6 替换 Endpoint.....	3427
5.13 ResourceQuota.....	3461
5.13.1 查询指定 namespace 下的 resourcequotas.....	3461
5.13.2 查询指定的 resourcequota.....	3476
5.14 CronJob.....	3487
5.14.1 查询所有 namespace 下的 CronJobs.....	3487
5.15 API groups.....	3565
5.15.1 查询 APIVersions.....	3565
5.15.2 查询所有 v1 版本的 API.....	3568
5.15.3 查询 APIGroupList.....	3572
5.15.4 查询 APIGroup /apis/apps.....	3577
5.15.5 查询所有 apps/v1 的 API.....	3581
5.15.6 查询 APIGroup /apis/batch.....	3585
5.15.7 查询 APIGroup /apis/batch.volcano.sh.....	3588
5.15.8 查询所有 batch.volcano.sh/v1alpha1 的 API.....	3591
5.15.9 查询所有 batch/v1 的 API.....	3594
5.15.10 查询所有 batch/v1beta1 的 API.....	3598
5.15.11 查询 APIGroup /apis/crd.yangtse.cni.....	3601
5.15.12 查询所有 crd.yangtse.cni/v1 的 API.....	3604
5.15.13 查询 APIGroup /apis/extensions.....	3608
5.15.14 查询所有 extensions/v1beta1 的 API.....	3611
5.15.15 查询 APIGroup /apis/metrics.k8s.io.....	3614
5.15.16 查询所有 metrics.k8s.io/v1beta1 的 API.....	3618
5.15.17 查询 APIGroup /apis/networking.cci.io.....	3621
5.15.18 查询所有 networking.cci.io/v1beta1 的 API.....	3625
5.15.19 查询 APIGroup /apis/rbac.authorization.k8s.io.....	3628
5.15.20 查询所有 rbac.authorization.k8s.io/v1 的 API.....	3632
5.16 Event.....	3635
5.16.1 查询指定 namespace 下的 Events.....	3635
5.16.2 删除 Event.....	3652
5.16.3 查询 Event.....	3664
5.17 PersistentVolumeClaim.....	3675
5.17.1 查询指定 namespace 下的 PersistentVolumeClaims.....	3675
5.17.2 创建 PersistentVolumeClaim.....	3692
5.17.3 删除 PersistentVolumeClaim.....	3740
5.17.4 查询 PersistentVolumeClaim.....	3766
5.17.5 替换 PersistentVolumeClaims.....	3779
5.18 RoleBinding.....	3817
5.18.1 获取指定 namespace 下 RoleBinding 列表.....	3817
5.18.2 创建 RoleBinding.....	3831

5.18.3 删除指定的 RoleBinding.....	3866
5.18.4 获取指定的 RoleBinding.....	3877
5.18.5 更新指定的 RoleBinding.....	3886
5.18.6 替换指定的 RoleBinding.....	3898
5.18.7 获取 RoleBinding 列表.....	3925
5.19 StatefulSet.....	3939
5.19.1 删除指定 namespace 下的 StatefulSets.....	3939
5.19.2 查询指定 namespace 下的 StatefulSets.....	3951
5.19.3 创建 StatefulSet.....	4031
5.19.4 删除 StatefulSet.....	4347
5.19.5 查询 StatefulSet.....	4359
5.19.6 更新 StatefulSet.....	4433
5.19.7 替换 StatefulSet.....	4509
5.19.8 查询 StatefulSet 状态.....	4754
5.19.9 查询用户所有的 StatefulSets.....	4828
5.20 Job.....	4908
5.20.1 查询用户所有的 Jobs.....	4908
5.20.2 删除指定 namespace 下的 Jobs.....	4986
5.20.3 查询指定 namespace 下的 Jobs.....	4998
5.20.4 创建 Job.....	5076
5.20.5 删除 Job.....	5389
5.20.6 查询 Job.....	5402
5.20.7 更新 Job.....	5475
5.20.8 替换 Job.....	5549
5.20.9 查询 Job 状态.....	5787
5.21 ReplicaSet.....	5859
5.21.1 查询指定 namespace 下的 ReplicaSets.....	5859
5.21.2 查询 ReplicaSet.....	5936
6 数据结构.....	6009
6.1 请求数据结构（废弃）.....	6009
6.2 响应数据结构（废弃）.....	6056
6.3 数据结构.....	6091
7 权限和授权项.....	6233
7.1 权限及授权项说明.....	6233
7.2 授权项分类.....	6235
8 附录.....	6250
8.1 PATCH 请求方法操作说明.....	6250
8.2 PATCH 请求方法操作示例.....	6250
8.3 约束限制.....	6254
8.4 Namespace 和 Network.....	6255
8.5 状态码.....	6260

8.6 错误码.....	6263
8.7 获取项目 ID.....	6283
8.8 获取账号 ID.....	6284
8.9 获取容器镜像地址.....	6285
9 历史 API.....	6286
9.1 Kubernetes API (废弃)	6286
9.1.1 Extended PersistentVolumeClaim.....	6286
9.1.1.1 导入存储.....	6286
9.1.1.2 查询导入的 PVC.....	6289
9.1.1.3 解绑存储.....	6292
9.1.2 TFJob.....	6294
9.1.2.1 创建 TFJob.....	6294
9.1.2.2 查询 TFJob.....	6298
9.1.2.3 查询指定 namespace 下的所有 TFJob.....	6301
9.1.2.4 删除 namespace 下的所有 TFJob.....	6305
9.1.2.5 删除 TFJob.....	6310
9.1.2.6 更新 TFJob.....	6312
9.1.2.7 替换 TFJob.....	6315
9.1.3 MXJob.....	6321
9.1.3.1 创建 MXJob.....	6321
9.1.3.2 查询 MXJob.....	6326
9.1.3.3 查询指定 namespace 下的所有 MXJob.....	6329
9.1.3.4 删除 namespace 下的所有 MXJob.....	6334
9.1.3.5 删除 MXJob.....	6339
9.1.3.6 更新 MXJob.....	6341
9.1.3.7 替换 MXJob.....	6344
9.1.4 PyTorchJob.....	6350
9.1.4.1 创建 PyTorchJob.....	6351
9.1.4.2 查询 PyTorchJob.....	6354
9.1.4.3 查询指定 namespace 下的所有 PyTorchJob.....	6357
9.1.4.4 删除 namespace 下的所有 PyTorchJob.....	6361
9.1.4.5 删除 PyTorchJob.....	6365
9.1.4.6 更新 PyTorchJob.....	6367
9.1.4.7 替换 PyTorchJob.....	6370
A 修订记录.....	6375

1 使用前必读

1.1 概述

欢迎使用云容器实例（Cloud Container Instance）。云容器实例提供 Serverless Container（无服务器容器）引擎，让您无需创建和管理服务器集群即可直接运行容器。

您可以使用本文档提供API对云容器实例进行相关操作，如创建、删除、变更规格等。

在调用云容器实例API之前，请确保已经充分了解云容器实例相关概念，并对 Kubernetes基本概念与知识有一定了解，详细信息请参见[产品介绍](#)。

1.2 调用说明

云容器实例提供了REST（Representational State Transfer）风格API，支持您通过HTTPS请求调用，调用方法请参见[如何调用API](#)。

1.3 终端节点

终端节点（Endpoint）即调用API的[请求地址](#)，不同服务不同区域的终端节点不同，您可以从[地区和终端节点](#)中查询所有服务的终端节点。

1.4 约束与限制

- 您能创建的云容器实例资源的数量与配额有关系，具体请参见[服务配额](#)。如果您需要扩大配额，请参见[如何申请扩大配额](#)。
- 更详细的限制请参见具体API的说明。

1.5 基本概念

- 账号

用户注册时的账号，账号对其所拥有的资源及云服务具有完全的访问权限，可以重置用户密码、分配用户权限等。由于账号是付费主体，为了确保账号安全，建议您不要直接使用账号进行日常管理工作，而是创建用户并使用用户进行日常管理工作。

- 用户

由账号在IAM中创建的用户，是云服务的使用人员，具有身份凭证（密码和访问密钥）。

通常在调用API的鉴权过程中，您需要用到账号、用户和密码等信息。

- 区域（Region）

从地理位置和网络时延维度划分，同一个Region内共享弹性计算、块存储、对象存储、VPC网络、弹性公网IP、镜像等公共服务。Region分为通用Region和专属Region，通用Region指面向公共租户提供通用云服务的Region；专属Region指只承载同一类业务或只面向特定租户提供业务服务的专用Region。

详情请参见[区域和可用区](#)。


- 可用区（AZ，Availability Zone）

一个可用区是一个或多个物理数据中心的集合，有独立的风火水电，AZ内逻辑上再将计算、网络、存储等资源划分成多个集群。一个Region中的多个AZ间通过高速光纤相连，以满足用户跨AZ构建高可用性系统的需求。

- 项目

区域默认对应一个项目，这个项目由系统预置，用来隔离物理区域间的资源（计算资源、存储资源和网络资源），以默认项目为单位进行授权，用户可以访问您账号中该区域的所有资源。如果您希望进行更加精细的权限控制，可以在区域默认的项目中创建子项目，并在子项目中创建资源，然后以子项目为单位进行授权，使得用户仅能访问特定子项目中的资源，使得资源的权限控制更加精确。

图 1-1 项目隔离模型



- 企业项目

企业项目是项目的升级版，针对企业不同项目间的资源进行分组和管理，是逻辑隔离。企业项目中可以包含多个区域的资源，且项目中的资源可以迁入迁出。

关于企业项目ID的获取及企业项目特性的详细信息，请参见《[企业管理用户指南](#)》。

1.6 API 版本选择建议

云容器实例提供了[Kubernetes API](#)，此版本跟随Kubernetes社区最新版本，且提供了老版本[Kubernetes API（OLD VERSIONS）](#)。

建议您使用Kubernetes API，能够更好的满足您的需求。

2 如何调用 API

2.1 构造请求

本节介绍REST API请求的组成，并以调用IAM服务的[管理员创建IAM用户](#)来说明如何调用API，该API获取用户的Token，Token可以用于调用其他API时鉴权。

请求 URI

请求URI由如下部分组成：

{URI-scheme}://{Endpoint}/{resource-path}?{query-string}

尽管请求URI包含在请求消息头中，但大多数语言或框架都要求您从请求消息中单独传递它，所以在此单独强调。

表 2-1 URI 中的参数说明

参数	描述
URI-scheme	表示用于传输请求的协议，当前所有API均采用HTTPS协议。
Endpoint	指定承载REST服务端点的服务器域名或IP，不同服务不同区域的Endpoint不同，您可以从 地区和终端节点 获取。 例如IAM服务在“中国-香港”区域的Endpoint为“iam.ap-southeast-1.myhuaweicloud.com”。
resource-path	资源路径，即API访问路径。从具体API的URI模块获取，例如“获取用户Token”API的resource-path为“/v3/auth/tokens”。
query-string	查询参数，是可选部分，并不是每个API都有查询参数。查询参数前面需要带一个“？”，形式为“参数名=参数取值”，例如“？limit=10”，表示查询不超过10条数据。

例如您需要获取IAM在“中国-香港”区域的Token，则需使用“中国-香港”区域的Endpoint（iam.ap-southeast-1.myhuaweicloud.com），并在[管理员创建IAM用户](#)的URI部分找到resource-path（/v3.0/OS-USER/users），拼接起来如下所示。

```
https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
```

图 2-1 URI 示意图



说明

为方便查看，在每个具体API的URI部分，只给出resource-path部分，并将请求方法写在一起。这是因为URI-scheme都是HTTPS，而Endpoint在同一个区域也相同，所以简洁起见将这两部分省略。

请求方法

HTTP请求方法（也称为操作或动词），它告诉服务您正在请求什么类型的操作。

表 2-2 HTTP 方法

方法	说明
GET	请求服务器返回指定资源。
PUT	请求服务器更新指定资源。
POST	请求服务器新增资源或执行特殊操作。
DELETE	请求服务器删除指定资源，如删除对象等。
HEAD	请求服务器资源头部。
PATCH	请求服务器更新资源的部分内容。 当资源不存在的时候，PATCH可能会去创建一个新的资源。

在[管理员创建IAM用户](#)的URI部分，您可以看到其请求方法为“POST”，则其请求为：

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
```

请求消息头

附加请求头字段，如指定的URI和HTTP方法所要求的字段。例如定义消息体类型的请求头“Content-Type”，请求鉴权信息等。

详细的公共请求消息头字段请参见[表2-3](#)。

表 2-3 公共请求消息头

名称	描述	是否必选	示例
Host	请求的服务器信息，从服务API的URL中获取。值为hostname[:port]。端口缺省时使用默认的端口，https的默认端口为443。	否 使用AK/SK认证时该字段必选。	code.test.com or code.test.com:443
Content-Type	消息体的类型（格式）。推荐用户使用默认值application/json，有其他取值时会在具体接口中专门说明。	是	application/json
Content-Length	请求body长度，单位为Byte。	否	3495
X-Project-Id	project id，项目编号。请参考 获取项目ID 章节获取项目编号。	否 如果是专属云场景采用AK/SK认证方式的接口请求，或者多project场景采用AK/SK认证的接口请求，则该字段必选。	e9993fc787d94b6c886cb aa340f9c0f4
X-Auth-Token	用户Token。 用户Token也就是调用 获取用户Token 接口的响应值，该接口是唯一不需要认证的接口。 请求响应成功后在响应消息头（Headers）中包含的“X-Subject-Token”的值即为Token值。	否 使用Token认证时该字段必选。	注：以下仅为Token示例片段。 MIIPAgYJKoZlhvcNAQcCo ...ggg1BBIIlNPXsidG9rZ

说明

API同时支持使用AK/SK认证，AK/SK认证使用SDK对请求进行签名，签名过程会自动往请求中添加Authorization（签名认证信息）和X-Sdk-Date（请求发送的时间）请求头。

AK/SK认证的详细说明请参见[认证鉴权](#)的“AK/SK认证”。

对于[管理员创建IAM用户](#)接口，使用AK/SK方式认证时，添加消息头后的请求如下所示。

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
Content-Type: application/json
X-Sdk-Date: 20240416T095341Z
Authorization: SDK-HMAC-SHA256 Access=*****, SignedHeaders=content-type;host;x-sdk-date,
Signature=*****
```

请求消息体（可选）

该部分可选。请求消息体通常以结构化格式（如JSON或XML）发出，与请求消息头中Content-Type对应，传递除请求消息头之外的内容。若请求消息体中的参数支持中文，则中文字符必须为UTF-8编码。

每个接口的请求消息体内容不同，也并不是每个接口都需要有请求消息体（或者说消息体为空），GET、DELETE操作类型的接口就不需要消息体，消息体具体内容需要根据具体接口而定。

对于[管理员创建IAM用户](#)接口，您可以从接口的请求部分看到所需的请求参数及参数说明，将消息体加入后的请求如下所示，其中加粗的字段需要根据实际值填写。

- **accountid**为IAM用户所属的账号ID。
- **username**为要创建的IAM用户名。
- **email**为IAM用户的邮箱。
- *********为IAM用户的登录密码。

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3.0/OS-USER/users
Content-Type: application/json
X-Sdk-Date: 20240416T095341Z
Authorization: SDK-HMAC-SHA256 Access=*****, SignedHeaders=content-type;host;x-sdk-date,
Signature=*****

{
  "user": {
    "domain_id": "accountid",
    "name": "username",
    "password": "*****",
    "email": "email",
    "description": "IAM User Description"
  }
}
```

到这里为止这个请求需要的内容就具备齐全了，您可以使用[curl](#)、[Postman](#)或直接编写代码等方式发送请求调用API。对于获取用户Token接口，返回的响应消息头中的“X-Subject-Token”就是需要获取的用户Token。有了Token之后，您就可以使用Token认证调用其他API。

2.2 认证鉴权

调用接口有如下两种认证方式，您可以选择其中一种进行认证鉴权。

- AK/SK认证：通过AK（Access Key ID）/SK（Secret Access Key）加密调用请求。推荐使用AK/SK认证，其安全性比Token认证要高。

- Token认证：通过Token认证调用请求。

AK/SK 认证

📖 说明

AK/SK签名认证方式仅支持消息体大小在12MB以内，12MB以上的请求请使用Token认证。

AK/SK认证就是使用AK/SK对请求进行签名，在请求时将签名信息添加到消息头，从而通过身份认证。

- AK (Access Key ID)：访问密钥ID。与私有访问密钥关联的唯一标识符；访问密钥ID和私有访问密钥一起使用，对请求进行加密签名。
- SK (Secret Access Key)：私有访问密钥。与访问密钥ID结合使用，对请求进行加密签名，可标识发送方，并防止请求被修改。

使用AK/SK认证时，您可以基于签名算法使用AK/SK对请求进行签名，也可以使用专门的签名SDK对请求进行签名。详细的签名方法和SDK使用方法请参见[API签名指南](#)。

📖 说明

签名SDK只提供签名功能，与服务提供的SDK不同，使用时请注意。

Token 认证

📖 说明

Token的有效期为24小时，需要使用一个Token鉴权时，可以先缓存起来，避免频繁调用。

Token在计算机系统中代表令牌（临时）的意思，拥有Token就代表拥有某种权限。Token认证就是在调用API的时候将Token加到请求消息头中，从而通过身份认证，获得操作API的权限。Token可通过调用[获取用户Token](#)接口获取。

调用本服务API需要项目级别的Token，即调用[获取用户Token](#)接口时，请求body中auth.scope的取值需要选择project，如下所示。

```
{
  "auth": {
    "identity": {
      "methods": [
        "password"
      ],
      "password": {
        "user": {
          "name": "username", //IAM用户名
          "password": $ADMIN_PASS, //IAM用户密码，建议在配置文件或者环境变量中密文存放，使用时解密，确保安全
          "domain": {
            "name": "domainname" //IAM用户所属账号名
          }
        }
      }
    },
    "scope": {
      "project": {
        "name": "xxxxxxx" //项目名称
      }
    }
  }
}
```


获取Token后，再调用其他接口时，您需要在请求消息头中添加“X-Auth-Token”，其值即为Token。例如Token值为“ABCDEFJ...”，则调用接口时将“X-Auth-Token: ABCDEFJ...”加到请求消息头即可，如下所示。

```
POST https://iam.ap-southeast-1.myhuaweicloud.com/v3/auth/projects
Content-Type: application/json
X-Auth-Token: ABCDEFJ....
```

2.3 返回结果

状态码

请求发送以后，您会收到响应，其中包含状态码、响应消息头和消息体。

状态码是一组从1xx到5xx的数字代码，状态码表示了请求响应的状态，完整的状态码列表请参见[状态码](#)。

对于[管理员创建IAM用户](#)接口，如果调用后返回状态码为“201”，则表示请求成功。

响应消息头

对应请求消息头，响应同样也有消息头，如“Content-type”。

对于[管理员创建IAM用户](#)接口，返回如[图2-2](#)所示的消息头，其中“X-Subject-Token”就是需要获取的用户Token。有了Token之后，您就可以使用Token认证调用其他API。

说明

建议在配置文件或者环境变量中密文存放，使用时解密，确保安全。

图 2-2 管理员创建 IAM 用户响应消息头

```
"X-Frame-Options": "SAMEORIGIN",
"X-IAM-ETag-id": "2562365939-d8f6f12921974cb097338ac11fceac8a",
"Transfer-Encoding": "chunked",
"Strict-Transport-Security": "max-age=31536000; includeSubdomains;",
"Server": "api-gateway",
"X-Request-Id": "af2953f2bcc67a42325a69a19e6c32a2",
"X-Content-Type-Options": "nosniff",
"Connection": "keep-alive",
"X-Download-Options": "noopen",
"X-XSS-Protection": "1; mode=block;",
"X-IAM-Trace-Id": "token_█_null_af2953f2bcc67a42325a69a19e6c32a2",
"Date": "Tue, 21 May 2024 09:03:40 GMT",
"Content-Type": "application/json; charset=utf8"
```

响应消息体（可选）

该部分可选。响应消息体通常以结构化格式（如JSON或XML）返回，与响应消息头中Content-Type对应，传递除响应消息头之外的内容。

对于[管理员创建IAM用户](#)接口，返回如下消息体。为篇幅起见，这里只展示部分内容。

```
{
  "user": {
```

```
"id": "c131886aec...",
"name": "IAMUser",
"description": "IAM User Description",
"areacode": "",
"phone": "",
"email": "****@***.com",
"status": null,
"enabled": true,
"pwd_status": false,
"access_mode": "default",
"is_domain_owner": false,
"xuser_id": "",
"xuser_type": "",
"password_expires_at": null,
"create_time": "2024-05-21T09:03:41.000000",
"domain_id": "d78cbac1.....",
"xdomain_id": "30086000.....",
"xdomain_type": "",
"default_project_id": null
}
}
```

当接口调用出错时，会返回错误码及错误信息说明，错误响应的Body体格式如下所示。

```
{
  "error_msg": "The request message format is invalid.",
  "error_code": "IMG.0001"
}
```

其中，error_code表示错误码，error_msg表示错误描述信息。

3 快速入门

概述

本文通过调用一系列云容器实例的API使用nginx镜像创建一个工作负载，介绍使用云容器实例API的基本流程。

其中，镜像直接使用开源镜像中心的nginx镜像，容器的资源规格为0.25核CPU、512M内存，并绑定一个负载均衡实例，通过负载均衡实例从外网访问容器负载。

API的调用方法请参见[如何调用API](#)。

创建流程

1. 调用[创建Namespace](#)接口创建命名空间。
2. 调用[创建Network](#)接口创建网络，与VPC和子网关关联。
3. 调用[创建Deployment](#)接口创建nginx容器负载。

创建 Nginx 负载

步骤1 调用[创建Namespace](#)接口创建命名空间，并指定使用GPU型资源。

```
{
  "apiVersion": "v1",
  "kind": "Namespace",
  "metadata": {
    "name": "namespace-test",
    "annotations": {
      "namespace.kubernetes.io/flavor": "gpu-accelerated"
    }
  },
  "spec": {
    "finalizers": [
      "kubernetes"
    ]
  }
}
```

您需要指定如下参数。

- name: 命名空间的名称。
- namespace.kubernetes.io/flavor: 命名空间的类型，当前支持GPU加速型（**gpu-accelerated**）和通用计算型（**general-computing**）。GPU加速型命名空间中可以使用GPU显卡。

步骤2 调用**创建Network**接口创建网络，与VPC与子网关关联。

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "kind": "Network",
  "metadata": {
    "name": "test-network",
    "annotations": {
      "network.alpha.kubernetes.io/default-security-group": "security-group-id",
      "network.alpha.kubernetes.io/domain-id": "domain-id",
      "network.alpha.kubernetes.io/project-id": "project-id"
    }
  },
  "spec": {
    "cidr": "192.168.0.0/24",
    "attachedVPC": "vpc-id",
    "networkID": "network-id",
    "subnetID": "subnet-id",
    "networkType": "underlay_neutron"
  }
}
```

您需要指定如下参数。

- name：网络对象的名称。
- network.alpha.kubernetes.io/default-security-group：安全组ID，您可以在[安全组控制台](#)查看。
- network.alpha.kubernetes.io/domain-id：账号ID，获取方法请参见[获取账号ID](#)。
- network.alpha.kubernetes.io/project-id：项目ID，获取方法请参见[获取项目ID](#)。
- cidr：网段地址，虚拟私有云下可用子网的范围。
- attachedVPC：命名空间所在VPC（虚拟私有云）的ID，您可以在[虚拟私有云控制台](#)中查询，也可以通过[查询VPC列表](#)API查询。
- networkID：VPC内子网的网络ID，您可以在[虚拟私有云控制台](#)中查询，也可以通过[查询子网列表](#)API查询。
- subnetID：VPC内子网的ID，您可以在[虚拟私有云控制台](#)中查询，也可以通过[查询子网列表](#)API查询。

步骤3 调用**创建Deployment**接口创建nginx容器负载。

负载名称为nginx，使用开源镜像中心的nginx:latest镜像，容器的资源规格为0.25核CPU、1G内存。调用接口后，云容器实例会创建一个运行nginx的容器。

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "name": "nginx"
  },
  "spec": {
    "replicas": 1,
    "selector": {
      "matchLabels": {
        "app": "nginx"
      }
    }
  },
  "template": {
    "metadata": {
      "labels": {
        "app": "nginx"
      }
    }
  }
}
```

```
    },  
    "spec": {  
      "containers": [  
        {  
          "image": "nginx:latest",  
          "name": "container-0",  
          "resources": {  
            "limits": {  
              "cpu": "250m",  
              "memory": "1Gi"  
            },  
            "requests": {  
              "cpu": "250m",  
              "memory": "1Gi"  
            }  
          }  
        }  
      ],  
      "imagePullSecrets": [  
        {  
          "name": "imagepull-secret"  
        }  
      ]  
    }  
  }  
}
```

您需要指定如下参数。

- name: Deployment的名称。
- replicas: Pod的数量，即Deployment下有几个Pod。
- selector.matchLabels: Deployment使用哪些标签选择Pod，例如这里设置的标签app=nginx，当Pod有这个标签时，就会被Deployment选中管理。
- template: Pod的模板，定义Pod的各种配置和规格。
 - metadata.labels: Pod的标签。
 - spec.containers: Pod中容器的定义。
 - image: 创建容器使用的镜像。
 - resources.limits: 容器使用资源的大小限制，即容器使用的资源不能超过这个限制。
 - resources.requests: 容器申请使用资源的大小。

Nginx负载创建完后，您可以在云容器实例控制台查看到负载。

图 3-1 nginx



---结束

4 API

4.1 Metrics

4.1.1 获取指定 namespace 下所有 pods 的 metrics 信息

功能介绍

获取指定namespace下所有pods的metrics信息

调用方法

请参见[如何调用API](#)。

URI

GET /apis/metrics.k8s.io/v1beta1/namespaces/{namespace}/pods

表 4-1 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-2 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 4-3 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-4 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
items	Array of io.k8s.metrics.pkg.apis.metrics.v1beta1.PodMetrics objects	List of pod metrics.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta_v2 object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds

表 4-5 io.k8s.metrics.pkg.apis.metrics.v1beta1.PodMetrics

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources

参数	参数类型	描述
containers	Array of io.k8s.metrics.pkg.apis.metrics.v1beta1.ContainerMetrics objects	Metrics for all containers are collected within the same time window.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta_v2 object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
timestamp	String	The following fields define time interval from which metrics were collected from the interval [Timestamp-Window, Timestamp].
window	String	Duration is a wrapper around time.Duration which supports correct marshaling to YAML and JSON. In particular, it marshals into strings, which can be used as map keys in json.

表 4-6 io.k8s.metrics.pkg.apis.metrics.v1beta1.ContainerMetrics

参数	参数类型	描述
name	String	Container name corresponding to the one from pod.spec.containers.
usage	Map<String,String>	The memory usage is the memory working set.

表 4-7 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta_v2

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata</p>
finalizers	Array of strings	<p>Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed.</p>

参数	参数类型	描述
generateName	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#idempotency</p>
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.
initializers	io.k8s.apimachinery.pkg.apis.meta.v1.Initializers object	<p>An initializer is a controller which enforces some system invariant at object creation time. This field is a list of initializers that have not yet acted on this object. If nil or empty, this object has been completely initialized. Otherwise, the object is considered uninitialized and is hidden (in list/watch and get calls) from clients that haven't explicitly asked to observe uninitialized objects.</p> <p>When an object is created, the system will populate this list with the current set of initializers. Only privileged users may set or modify this list. Once it is empty, it may not be modified further by any user.</p> <p>DEPRECATED - initializers are an alpha field and will be removed in v1.15.</p>
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels

参数	参数类型	描述
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry_v2 objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object. This field is alpha and can be changed or removed without notice.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference_v2 objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-8 io.k8s.apimachinery.pkg.apis.meta.v1.Initializers

参数	参数类型	描述
pending	Array of io.k8s.apimachinery.pkg.apis.meta.v1.Initializer objects	Pending is a list of initializers that must execute in order before this object is visible. When the last pending initializer is removed, and no failing result is set, the initializers struct will be set to nil and the object is considered as initialized and visible to all clients.
result	io.k8s.apimachinery.pkg.apis.meta.v1.Status_v2 object	If result is set with the Failure field, the object will be persisted to storage and then deleted, ensuring that other clients can observe the deletion.

表 4-9 io.k8s.apimachinery.pkg.apis.meta.v1.Initializer

参数	参数类型	描述
name	String	name of the process that is responsible for initializing this object.

表 4-10 io.k8s.apimachinery.pkg.apis.meta.v1.Status_v2

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails_v2 object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta_v2 object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.

参数	参数类型	描述
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#spec-and-status

表 4-11 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails_v2

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-12 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 4-13 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry_v2

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fields	Object	Fields identifies a set of fields.
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-14 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference_v2

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-15 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta_v2

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "metrics.k8s.io/v1beta1",
  "items": [ {
    "containers": [ {
      "name": "container-0",
      "usage": {
        "cpu": "374598n",
        "cpu_accumulated": "127739654828n",
        "memory": "1540Ki",
        "memory_rss": "284Ki",
        "memory_usage": "1804Ki"
      }
    }
  ]},
  "metadata": {
    "creationTimestamp": "2023-03-28T11:24:34Z",
    "name": "cci-deployment-20233231-6c8d7f74f7-2t6k5",
    "namespace": "cci-namespace-13238101",
    "selfLink": "/apis/metrics.k8s.io/v1beta1/namespaces/cci-namespace-13238101/pods/cci-deployment-20233231-6c8d7f74f7-2t6k5"
  },
  "timestamp": "2023-03-28T11:24:19Z",
  "window": "30s"
}, {
  "containers": [ {
    "name": "container-0",
    "usage": {
      "cpu": "0",
      "cpu_accumulated": "75422933n",
      "memory": "6236Ki",
      "memory_rss": "1644Ki",
      "memory_usage": "12308Ki"
    }
  }
  ]},
  "metadata": {
    "creationTimestamp": "2023-03-28T11:24:34Z",
```

```
"name": "cci-deployment-687-1678883958925-599b74469b-7bmqh",
"namespace": "cci-namespace-13238101",
"selfLink": "/apis/metrics.k8s.io/v1beta1/namespaces/cci-namespace-13238101/pods/cci-
deployment-687-1678883958925-599b74469b-7bmqh"
},
"timestamp": "2023-03-28T11:24:15Z",
"window": "30s"
}],
"kind": "PodMetricsList",
"metadata": {
  "selfLink": "/apis/metrics.k8s.io/v1beta1/namespaces/cci-namespace-13238101/pods"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.1.2 获取指定 namespace 下指定 pod 的 metrics 信息

功能介绍

获取指定namespace下指定pod的metrics信息

调用方法

请参见[如何调用API](#)。

URI

GET /apis/metrics.k8s.io/v1beta1/namespaces/{namespace}/pods/{name}

表 4-16 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the PodMetrics
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-17 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-18 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-19 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources

参数	参数类型	描述
containers	Array of io.k8s.metrics.pkg.apis.metrics.v1beta1.ContainerMetrics objects	Metrics for all containers are collected within the same time window.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta_v2 object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
timestamp	String	The following fields define time interval from which metrics were collected from the interval [Timestamp-Window, Timestamp].
window	String	Duration is a wrapper around time.Duration which supports correct marshaling to YAML and JSON. In particular, it marshals into strings, which can be used as map keys in json.

表 4-20 io.k8s.metrics.pkg.apis.metrics.v1beta1.ContainerMetrics

参数	参数类型	描述
name	String	Container name corresponding to the one from pod.spec.containers.
usage	Map<String,String>	The memory usage is the memory working set.

表 4-21 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta_v2

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata</p>
finalizers	Array of strings	<p>Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed.</p>

参数	参数类型	描述
generateName	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#idempotency</p>
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.
initializers	io.k8s.apimachinery.pkg.apis.meta.v1.Initializers object	<p>An initializer is a controller which enforces some system invariant at object creation time. This field is a list of initializers that have not yet acted on this object. If nil or empty, this object has been completely initialized. Otherwise, the object is considered uninitialized and is hidden (in list/watch and get calls) from clients that haven't explicitly asked to observe uninitialized objects.</p> <p>When an object is created, the system will populate this list with the current set of initializers. Only privileged users may set or modify this list. Once it is empty, it may not be modified further by any user.</p> <p>DEPRECATED - initializers are an alpha field and will be removed in v1.15.</p>
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels

参数	参数类型	描述
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry_v2 objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object. This field is alpha and can be changed or removed without notice.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference_v2 objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-22 io.k8s.apimachinery.pkg.apis.meta.v1.Initializers

参数	参数类型	描述
pending	Array of io.k8s.apimachinery.pkg.apis.meta.v1.Initializer objects	Pending is a list of initializers that must execute in order before this object is visible. When the last pending initializer is removed, and no failing result is set, the initializers struct will be set to nil and the object is considered as initialized and visible to all clients.
result	io.k8s.apimachinery.pkg.apis.meta.v1.Status_v2 object	If result is set with the Failure field, the object will be persisted to storage and then deleted, ensuring that other clients can observe the deletion.

表 4-23 io.k8s.apimachinery.pkg.apis.meta.v1.Initializer

参数	参数类型	描述
name	String	name of the process that is responsible for initializing this object.

表 4-24 io.k8s.apimachinery.pkg.apis.meta.v1.Status_v2

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails_v2 object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta_v2 object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.

参数	参数类型	描述
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#spec-and-status

表 4-25 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails_v2

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-26 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 4-27 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta_v2

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#concurrency-control-and-consistency

参数	参数类型	描述
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only.

表 4-28 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry_v2

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fields	Object	Fields identifies a set of fields.
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-29 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference_v2

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds

参数	参数类型	描述
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "metrics.k8s.io/v1beta1",
  "containers": [ {
    "name": "container-0",
    "usage": {
      "cpu": "353450n",
      "cpu_accumulated": "127810375798n",
      "memory": "1560Ki",
      "memory_rss": "284Ki",
      "memory_usage": "1824Ki"
    }
  } ],
  "kind": "PodMetrics",
  "metadata": {
    "creationTimestamp": "2023-03-28T11:28:23Z",
    "name": "cci-deployment-20233231-6c8d7f74f7-2t6k5",
    "namespace": "cci-namespace-13238101",
    "selfLink": "/apis/metrics.k8s.io/v1beta1/namespaces/cci-namespace-13238101/pods/cci-deployment-20233231-6c8d7f74f7-2t6k5"
  },
  "timestamp": "2023-03-28T11:28:02Z",
  "window": "30s"
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2 EIPPool

4.2.1 查询指定 namespace 下的 EIP Pools

功能介绍

查询指定 namespace 下所有 EIP Pool 的详细信息。

调用方法

请参见[如何调用 API](#)。

URI

GET /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools

表 4-30 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-31 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-32 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-33 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of cni.yangtse.crd.v1.EIPPool objects	List of eippools. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 4-34 cni.yangtse.crd.v1.EIPPool

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-35 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-36 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-37 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-38 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-39 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-40 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值: 0 最大值: 200000

表 4-41 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-42 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-43 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-44 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-45 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceID	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

表 4-46 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "items": {
    "apiVersion": "crd.yangtse.cni/v1",
    "kind": "EIPPool",
    "metadata": {
      "creationTimestamp": "2022-09-07T01:22:50Z",
      "finalizers": [ "yangtse.io/eip-pool" ],
      "generation": 1,
      "name": "eippool-test",
      "namespace": "namespace-test",
      "resourceVersion": "42396258",
      "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
      "uid": "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
    },
    "spec": {
      "amount": 1,
      "eipAttributes": {
        "bandwidth": {
          "chargeMode": "bandwidth",
          "name": "eip-test",
          "shareType": "PER",
          "size": 5
        },
        "ipVersion": 4,
        "networkType": "5_g-vm"
      }
    },
    "status": {
      "eips": [ {
        "alias": "eip-test",
        "bandWidthChargeMode": "bandwidth",
        "bandwidthShareType": "PER",
        "bandwidthSize": 5,
        "id": "034a0bae-81f7-46f4-b933-3273adc32b54",
        "ipv4": "100.85.221.2",
        "networkType": "5_g-vm",
        "status": "DOWN"
      } ],
      "usage": "0/1"
    }
  },
  "kind": "EIPPoolList",
  "metadata": {
    "continue": "",
    "resourceVersion": "42396272",
    "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools"
  }
}
```

状态码

状态码	描述
200	OK

状态码	描述
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2.2 创建 EIPPool

功能介绍

创建EIPPool。

调用方法

请参见[如何调用API](#)。

URI

POST /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools

表 4-47 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-48 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-49 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 4-50 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	spec object	Specification of the EIPPool.
status	否	status object	Status of the EIPPool.

表 4-51 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	是否必选	参数类型	描述
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-52 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-53 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-54 spec

参数	是否必选	参数类型	描述
amount	否	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500
eipAttributes	否	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	否	Array of strings	EIPs used to generate EIP resources.

表 4-55 cni.yangtse.crd.v1.EIPAttributes

参数	是否必选	参数类型	描述
alias	否	String	Alias of PublicIP
bandwidth	是	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	否	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	是	String	Network Type of PublicIP

表 4-56 cni.yangtse.crd.v1.EIPBandwidth

参数	是否必选	参数类型	描述
chargeMode	否	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	否	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	否	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	否	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE
size	否	Integer	Bandwidth size 最小值：0 最大值：200000

表 4-57 status

参数	是否必选	参数类型	描述
associates	否	Map<String,c ni.yangtse.crd.v1.EIPAssociate>	eip associate infos.
eips	否	Array of eips objects	EIPs is a set of generated EIP resources.
usage	否	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-58 cni.yangtse.crd.v1.EIPAssociate

参数	是否必选	参数类型	描述
attachment	否	attachment object	Attachment is the eip attachment info.
privateIP	否	privateIP object	PrivateIP is private ip information

表 4-59 attachment

参数	是否必选	参数类型	描述
freeTimestamp	否	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	否	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	否	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	否	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	否	String	Pod name
podNamespace	否	String	Pod Namespace

表 4-60 privateIP

参数	是否必选	参数类型	描述
id	否	String	Port ID or EIP ID
ipv4	否	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	否	String	IPv6 is the IPv6 address for traffic from the eni.
mac	否	String	MAC is the MAC address of the endpoint interface.
status	否	String	Status of this IP

表 4-61 eips

参数	是否必选	参数类型	描述
alias	否	String	Alias of the PublicIP
associateInstanceID	否	String	Associate instance id
associateInstanceType	否	String	Associate instance type
bandwidthChargeMode	否	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	否	String	Bandwidth share type of the PublicIP
bandwidthSize	否	Integer	Bandwidth size of the PublicIP
id	否	String	The ID of the PublicIP
ipv4	否	String	The ipv4 address of the PublicIP
ipv6	否	String	The ipv6 address of the PublicIP
networkType	否	String	Network Type of PublicIP
status	否	String	PublicIP status

响应参数

状态码： 200

表 4-62 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-63 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-64 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-65 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-66 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-67 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-68 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值: 0 最大值: 200000

表 4-69 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-70 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-71 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-72 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-73 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

状态码： 201

表 4-74 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-75 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-76 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-77 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-78 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-79 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-80 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值: 0 最大值: 200000

表 4-81 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-82 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-83 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-84 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-85 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

状态码： 202

表 4-86 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-87 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-88 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-89 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-90 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-91 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-92 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值: 0 最大值: 200000

表 4-93 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-94 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-95 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-96 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-97 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

请求示例

动态创建独占带宽类型的EIPPool，资源池中EIP数量为2。

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "name": "eippool-test"
  },
  "spec": {
    "amount": 2,
    "eipAttributes": {
      "bandwidth": {
        "chargeMode": "bandwidth",
        "name": "eippool-test",
        "shareType": "PER",
        "size": 5
      },
      "ipVersion": 4,
      "networkType": "5_g-vm"
    }
  }
}
```

响应示例

状态码： 201

Created

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "creationTimestamp": "2023-04-23T06:36:06Z",
    "generation": 1,
    "name": "eippool-test",
    "namespace": "auto-test-namespace",
    "resourceVersion": "290279663",
    "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/auto-test-namespace/eippools/eippool-test",
    "uid": "446b63bb-c1f1-4eb2-8bdf-b77823881d37"
  },
  "spec": {
    "amount": 2,
    "eipAttributes": {
      "bandwidth": {
        "chargeMode": "bandwidth",
        "name": "eippool-test",
        "shareType": "PER",
        "size": 5
      },
      "ipVersion": 4,
      "networkType": "5_g-vm"
    }
  }
}
```

状态码

状态码	描述
200	OK
201	Created

状态码	描述
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2.3 删除指定的 EIPPool

功能介绍

删除指定的EIPPool。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools/{name}

表 4-98 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the EIPPool
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-99 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-100 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 4-101 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 4-102 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 4-103 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 4-104 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-105 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 4-106 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 4-107 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 4-108 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-109 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 4-110 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

删除指定的EIPPool

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions",
  "propagationPolicy": "Orphan"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "crd.yangtse.cni/v1",
```

```
"kind" : "EIPPool",
"metadata" : {
  "creationTimestamp" : "2022-09-07T01:22:50Z",
  "deletionGracePeriodSeconds" : 0,
  "deletionTimestamp" : "2022-09-07T01:22:59Z",
  "finalizers" : [ "yangtse.io/eip-pool" ],
  "generation" : 1,
  "name" : "eippool-test",
  "namespace" : "namespace-test",
  "resourceVersion" : "42396258",
  "selfLink" : "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
  "uid" : "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
},
"spec" : {
  "amount" : 1,
  "eipAttributes" : {
    "bandwidth" : {
      "chargeMode" : "bandwidth",
      "name" : "eip-test",
      "shareType" : "PER",
      "size" : 5
    },
    "ipVersion" : 4,
    "networkType" : "5_g-vm"
  }
},
"status" : {
  "eips" : [ {
    "alias" : "eip-test",
    "bandWidthChargeMode" : "bandwidth",
    "bandwidthShareType" : "PER",
    "bandwidthSize" : 5,
    "id" : "034a0bae-81f7-46f4-b933-3273adc32b54",
    "ipv4" : "100.85.221.2",
    "networkType" : "5_g-vm",
    "status" : "DOWN"
  } ],
  "usage" : "0/1"
}
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2.4 查询指定的 EIPPool

功能介绍

查询指定的EIPPool的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools/{name}

表 4-111 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the EIPPool
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-112 Query 参数

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-113 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-114 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-115 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-116 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-117 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-118 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-119 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-120 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值: 0 最大值: 200000

表 4-121 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-122 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-123 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-124 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-125 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "creationTimestamp": "2022-09-07T01:22:50Z",
    "finalizers": [ "yangtse.io/eip-pool" ],
    "generation": 1,
    "name": "eippool-test",
    "namespace": "namespace-test",
    "resourceVersion": "42396258",
    "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
    "uid": "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
  },
  "spec": {
    "amount": 1,
    "eipAttributes": {
      "bandwidth": {
        "chargeMode": "bandwidth",
        "name": "eip-test",
        "shareType": "PER",
        "size": 5
      },
      "ipVersion": 4,
      "networkType": "5_g-vm"
    }
  },
  "status": {
    "eips": [ {
      "alias": "eip-test",
      "bandwidthChargeMode": "bandwidth",
      "bandwidthShareType": "PER",
      "bandwidthSize": 5,
      "id": "034a0bae-81f7-46f4-b933-3273adc32b54",
      "ipv4": "100.85.221.2",
      "networkType": "5_g-vm",
      "status": "DOWN"
    } ],
    "usage": "0/1"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2.5 更新 EIPPool

功能介绍

更新EIPPool。

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools/{name}

表 4-126 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the EIPPool
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-127 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-128 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 4-129 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 4-130 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-131 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-132 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-133 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-134 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-135 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-136 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值：0 最大值：200000

表 4-137 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-138 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-139 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-140 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-141 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

请求示例

更新EIPPool中的labels值为"some-key":"some-value"。

```
{
  "metadata": {
    "labels": {
      "some-key": "some-value"
    }
  }
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "creationTimestamp": "2022-09-07T01:22:50Z",
    "finalizers": [ "yangtse.io/eip-pool" ],
    "generation": 1,
    "labels": {
      "some-key": "some-value"
    }
  },
  "name": "eippool-test",
  "namespace": "namespace-test",
  "resourceVersion": "42396258",
  "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
  "uid": "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
},
"spec": {
  "amount": 1,
  "eipAttributes": {
    "bandwidth": {
      "chargeMode": "bandwidth",
      "name": "eip-test",
      "shareType": "PER",
      "size": 5
    }
  },
  "ipVersion": 4,
  "networkType": "5_g-vm"
}
},
"status": {
  "eips": [ {
    "alias": "eip-test",
    "bandwidthChargeMode": "bandwidth",
    "bandwidthShareType": "PER",
    "bandwidthSize": 5,
    "id": "034a0bae-81f7-46f4-b933-3273adc32b54",
    "ipv4": "100.85.221.2",
    "networkType": "5_g-vm",
    "status": "DOWN"
  } ],
  "usage": "0/1"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2.6 替换 EIPPool

功能介绍

替换EIPPool。

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools/{name}

表 4-142 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the EIPPool
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-143 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-144 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none"> 1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。 2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。 3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。
Content-Type	是	String	<p>消息体的类型（格式），默认取值为“application/json” 缺省值：application/json</p>

表 4-145 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	spec object	Specification of the EIPPool.
status	否	status object	Status of the EIPPool.

表 4-146 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	是否必选	参数类型	描述
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-147 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-148 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-149 spec

参数	是否必选	参数类型	描述
amount	否	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500
eipAttributes	否	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	否	Array of strings	EIPs used to generate EIP resources.

表 4-150 cni.yangtse.crd.v1.EIPAttributes

参数	是否必选	参数类型	描述
alias	否	String	Alias of PublicIP
bandwidth	是	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	否	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	是	String	Network Type of PublicIP

表 4-151 cni.yangtse.crd.v1.EIPBandwidth

参数	是否必选	参数类型	描述
chargeMode	否	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	否	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	否	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	否	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE
size	否	Integer	Bandwidth size 最小值：0 最大值：200000

表 4-152 status

参数	是否必选	参数类型	描述
associates	否	Map<String,c ni.yangtse.crd.v1.EIPAssociate>	eip associate infos.
eips	否	Array of eips objects	EIPs is a set of generated EIP resources.
usage	否	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-153 cni.yangtse.crd.v1.EIPAssociate

参数	是否必选	参数类型	描述
attachment	否	attachment object	Attachment is the eip attachment info.
privateIP	否	privateIP object	PrivateIP is private ip information

表 4-154 attachment

参数	是否必选	参数类型	描述
freeTimestamp	否	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	否	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	否	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	否	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	否	String	Pod name
podNamespace	否	String	Pod Namespace

表 4-155 privateIP

参数	是否必选	参数类型	描述
id	否	String	Port ID or EIP ID
ipv4	否	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	否	String	IPv6 is the IPv6 address for traffic from the eni.
mac	否	String	MAC is the MAC address of the endpoint interface.
status	否	String	Status of this IP

表 4-156 eips

参数	是否必选	参数类型	描述
alias	否	String	Alias of the PublicIP
associateInstanceID	否	String	Associate instance id
associateInstanceType	否	String	Associate instance type
bandwidthChargeMode	否	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	否	String	Bandwidth share type of the PublicIP
bandwidthSize	否	Integer	Bandwidth size of the PublicIP
id	否	String	The ID of the PublicIP
ipv4	否	String	The ipv4 address of the PublicIP
ipv6	否	String	The ipv6 address of the PublicIP
networkType	否	String	Network Type of PublicIP
status	否	String	PublicIP status

响应参数

状态码： 200

表 4-157 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-158 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-159 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-160 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-161 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-162 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-163 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值：0 最大值：200000

表 4-164 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-165 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-166 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-167 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-168 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

状态码： 201

表 4-169 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-170 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-171 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-172 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-173 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-174 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-175 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值: 0 最大值: 200000

表 4-176 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-177 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-178 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-179 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-180 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

请求示例

将已创建EIPPool中的label值替换为"some-key" : "some-value"。

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "creationTimestamp": "2022-09-07T01:22:50Z",
    "finalizers": [ "yangtse.io/eip-pool" ],
    "generation": 1,
    "labels": {
      "some-key": "some-value"
    },
  },
  "name": "eippool-test",
  "namespace": "namespace-test",
  "resourceVersion": "42396258",
  "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
  "uid": "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
},
"spec": {
  "amount": 1,
  "eipAttributes": {
    "bandwidth": {
      "chargeMode": "bandwidth",
      "name": "eip-test",
      "shareType": "PER",
      "size": 5
    },
  },
  "ipVersion": 4,
  "networkType": "5_g-vm"
}
},
"status": {
  "eips": [ {
    "alias": "eip-test",
    "bandWidthChargeMode": "bandwidth",
    "bandwidthShareType": "PER",
    "bandwidthSize": 5,
    "id": "034a0bae-81f7-46f4-b933-3273adc32b54",
    "ipv4": "100.85.221.2",
    "networkType": "5_g-vm",
    "status": "DOWN"
  } ],
  "usage": "0/1"
}
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "creationTimestamp": "2022-09-07T01:22:50Z",
    "finalizers": [ "yangtse.io/eip-pool" ],
    "generation": 1,
    "labels": {
      "some-key": "some-value"
    },
  },
  "name": "eippool-test",
  "namespace": "namespace-test",
  "resourceVersion": "42396258",
  "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
}
```

```
{
  "uid": "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
},
"spec": {
  "amount": 1,
  "eipAttributes": {
    "bandwidth": {
      "chargeMode": "bandwidth",
      "name": "eip-test",
      "shareType": "PER",
      "size": 5
    },
    "ipVersion": 4,
    "networkType": "5_g-vm"
  }
},
"status": {
  "eips": [ {
    "alias": "eip-test",
    "bandWidthChargeMode": "bandwidth",
    "bandwidthShareType": "PER",
    "bandwidthSize": 5,
    "id": "034a0bae-81f7-46f4-b933-3273adc32b54",
    "ipv4": "100.85.221.2",
    "networkType": "5_g-vm",
    "status": "DOWN"
  } ],
  "usage": "0/1"
}
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.2.7 查询指定的 EIPPool

功能介绍

查询指定的EIPPool的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/crd.yangtse.cni/v1/namespaces/{namespace}/eippools/{name}/status

表 4-181 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the EIPPool
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-182 Query 参数

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-183 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-184 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	spec object	Specification of the EIPPool.
status	status object	Status of the EIPPool.

表 4-185 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-186 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-187 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-188 spec

参数	参数类型	描述
amount	Integer	Amount is the amount of eips need to be create. 最小值: 0 最大值: 500

参数	参数类型	描述
eipAttributes	cni.yangtse.crd.v1.EIPAttributes object	Attributes is the eip attributes which used to create eip.
eips	Array of strings	EIPs used to generate EIP resources.

表 4-189 cni.yangtse.crd.v1.EIPAttributes

参数	参数类型	描述
alias	String	Alias of PublicIP
bandwidth	cni.yangtse.crd.v1.EIPBandwidth object	Bandwidth create attributes
ipVersion	Integer	IP Version of PublicIP 枚举值： <ul style="list-style-type: none">• 4• 6
networkType	String	Network Type of PublicIP

表 4-190 cni.yangtse.crd.v1.EIPBandwidth

参数	参数类型	描述
chargeMode	String	Bandwidth Charge Mode 枚举值： <ul style="list-style-type: none">• bandwidth• traffic
id	String	Bandwidth id, only useful for WHOLE bandwidth share type
name	String	Bandwidth name, only useful for PER bandwidth share type 最小长度：1 最大长度：64
shareType	String	Bandwidth Share Type 枚举值： <ul style="list-style-type: none">• PER• WHOLE

参数	参数类型	描述
size	Integer	Bandwidth size 最小值：0 最大值：200000

表 4-191 status

参数	参数类型	描述
associates	Map<String, cni.yangtse.crd.v1.EIPAssociate >	eip associate infos.
eips	Array of eips objects	EIPs is a set of generated EIP resources.
usage	String	Usage is usage of eip in this pool, e.g. 2/10.

表 4-192 cni.yangtse.crd.v1.EIPAssociate

参数	参数类型	描述
attachment	attachment object	Attachment is the eip attachment info.
privateIP	privateIP object	PrivateIP is private ip information

表 4-193 attachment

参数	参数类型	描述
freeTimestamp	String	Timestamp of NeutronPort be to FixedPreBound or FixedUnBound phase
neutronPortName	String	NeutronPortName is NeutronPort name which associated.
neutronPortNamespace	String	NeutronPortNamespace is NeutronPort namespace which associated.
nodeName	String	nodeName is node name where the NeutronPort resides, must with phase Bound, PreBound or FixedBound.
podName	String	Pod name

参数	参数类型	描述
podNamespace	String	Pod Namespace

表 4-194 privateIP

参数	参数类型	描述
id	String	Port ID or EIP ID
ipv4	String	IPv4 is the IPv4 address for traffic from the eni.
ipv6	String	IPv6 is the IPv6 address for traffic from the eni.
mac	String	MAC is the MAC address of the endpoint interface.
status	String	Status of this IP

表 4-195 eips

参数	参数类型	描述
alias	String	Alias of the PublicIP
associateInstanceId	String	Associate instance id
associateInstanceType	String	Associate instance type
bandwidthChargeMode	String	Bandwidth charge mode of the PublicIP
bandwidthShareType	String	Bandwidth share type of the PublicIP
bandwidthSize	Integer	Bandwidth size of the PublicIP
id	String	The ID of the PublicIP
ipv4	String	The ipv4 address of the PublicIP
ipv6	String	The ipv6 address of the PublicIP
networkType	String	Network Type of PublicIP
status	String	PublicIP status

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "crd.yangtse.cni/v1",
  "kind": "EIPPool",
  "metadata": {
    "creationTimestamp": "2022-09-07T01:22:50Z",
    "finalizers": [ "yangtse.io/eip-pool" ],
    "generation": 1,
    "name": "eippool-test",
    "namespace": "namespace-test",
    "resourceVersion": "42396258",
    "selfLink": "/apis/crd.yangtse.cni/v1/namespaces/namespace-test/eippools/eippool-test",
    "uid": "e4dc5432-1d9b-4fcb-8840-ee445b6511ae"
  },
  "spec": {
    "amount": 1,
    "eipAttributes": {
      "bandwidth": {
        "chargeMode": "bandwidth",
        "name": "eip-test",
        "shareType": "PER",
        "size": 5
      },
      "ipVersion": 4,
      "networkType": "5_g-vm"
    }
  },
  "status": {
    "eips": [ {
      "alias": "eip-test",
      "bandWidthChargeMode": "bandwidth",
      "bandwidthShareType": "PER",
      "bandwidthSize": 5,
      "id": "034a0bae-81f7-46f4-b933-3273adc32b54",
      "ipv4": "100.85.221.2",
      "networkType": "5_g-vm",
      "status": "DOWN"
    } ],
    "usage": "0/1"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.3 Network

4.3.1 删除指定 namespace 下的 Networks

功能介绍

删除指定namespace下的所有Network对象。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks

表 4-196 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-197 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-198 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 4-199 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 4-200 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 4-201 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 4-202 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none"> for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-203 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 4-204 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

删除指定namespace下的所有Network对象

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions",
  "propagationPolicy": "Orphan"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "code": 200,
  "kind": "Status",
  "metadata": { },
  "status": "Success"
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.3.2 查询指定 namespace 下的 Networks

功能介绍

查询指定namespace下的所有Network对象。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks

表 4-205 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-206 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-207 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-208 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.cci.v1beta1.Network objects	Network is a network resource in container.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata

表 4-209 io.k8s.api.cci.v1beta1.Network

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-210 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 4-211 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-212 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-213 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-214 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-215 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

表 4-216 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "items": [ {
    "metadata": {
      "annotations": {
        "network.alpha.kubernetes.io/default-security-group": "19c5d024-aed5-4856-b958-c0f65ce70855",
        "network.alpha.kubernetes.io/domain-id": "aadb43c0b14c4cafbccfff483d075987",
        "network.alpha.kubernetes.io/project-id": "51bf52609f2a49c68bfda3398817b376"
      },
      "creationTimestamp": "2018-09-03T11:21:00Z",
      "name": "namespace-test-dc1-default-network",
      "namespace": "namespace-test",
      "resourceVersion": "5016899",
      "selfLink": "/apis/networking.cci.io/v1beta1/namespaces/namespace-test/networks/namespace-test-dc1-default-network",
      "uid": "6fb85414-af6b-11e8-b6ef-f898ef6c78b4"
    },
    "spec": {
      "attachedVPC": "0d4080e5-546a-46c4-86fe-f3e26d685177",
      "availableZone": "cn-north-4a",
      "cidr": "192.168.244.0/23",
      "networkID": "0022e356-f730-4226-802e-9cdaa6e7da17",
      "networkType": "underlay_neutron",
      "subnetID": "1ffd839d-e534-4fa8-a59d-42356335bf74"
    },
    "status": {
      "state": "Active"
    }
  } ],
  "kind": "NetworkList",
  "metadata": {
    "resourceVersion": "5016953",
    "selfLink": "/apis/networking.cci.io/v1beta1/namespaces/namespace-test/networks"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable

状态码	描述
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.3.3 创建 Network

功能介绍

创建一个Network对象。

Network对象是华为云CCI新增对象，用于定义kubernetes中一个namespace内的网络。目前CCI支持VPC网络，一个VPC网络类型的network对象对应于华为云虚拟私有云服务中的一个子网。

CCI的容器网络依赖于华为云底层VPC网络，因此在创建network对象前，需要先调用虚拟私有云的接口创建或者查询已有子网信息。

📖 说明

须知：此处VPC和子网的网段不能为10.247.0.0/16，10.247.0.0/16是云容器实例预留给Service的网段。如果您使用此网段，后续可能会造成IP冲突，导致负载无法创建或服务不可用；如果您不需要通过Service访问，而是直接访问Pod，则可以使用此网段。

具体如何创建一个Network对象，可以参考[Namespace和Network](#)

调用方法

请参见[如何调用API](#)。

URI

POST /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks

表 4-217 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-218 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-219 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 4-220 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	否	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	否	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-221 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-222 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-223 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-224 io.k8s.api.cci.v1beta1.NetworkSpec

参数	是否必选	参数类型	描述
attachedVPC	否	String	ID of the VPC to attach
availableZone	否	String	available zone
cidr	否	String	The CIDR of the network
networkID	否	String	network ID
networkType	否	String	network type like 'overlay_l2'
subnetID	否	String	Subnet ID

参数	是否必选	参数类型	描述
subnets	否	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-225 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	是否必选	参数类型	描述
cidr	否	String	The CIDR of the network
networkID	否	String	NetworkID describes the subnet network id
subnetID	否	String	SubnetID describes the subnet id

表 4-226 io.k8s.api.cci.v1beta1.NetworkStatus

参数	是否必选	参数类型	描述
message	否	String	Message describes why network is in current state
state	否	String	State describes the network state

响应参数

状态码： 200

表 4-227 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-228 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 4-229 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-230 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-231 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-232 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-233 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

状态码： 201

表 4-234 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-235 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-236 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-237 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-238 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-239 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-240 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

状态码： 202

表 4-241 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-242 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-243 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-244 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-245 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-246 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-247 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

请求示例

创建Network,指定账号ID、项目ID和安全组ID, 关联VPC和子网网段。
metadata.annotations必选字段数据结构说明

参数	是否必选	参数类型	描述
network.alpha.kubernetes.io/default-security-group	Yes	String	Network对应子网所属安全组ID。
network.alpha.kubernetes.io/project-id	Yes	String	Network所属用户的project-id, 获取方法请参见 获取项目ID 。

参数	是否必选	参数类型	描述
network.alpha.kubernetes.io/domain-id	Yes	String	Network所属用户的domain-id，获取方法请参见获取账号ID。

VPC网络类型

参数	是否必选	参数类型	描述
cidr	No	String	Network对应VPC子网的网段。 **说明：** 此处VPC和子网的网段不能为10.247.0.0/16，10.247.0.0/16是云容器实例预留给Service的网段。如果您使用此网段，后续可能会造成IP冲突，导致负载无法创建或服务不可用；如果您不需要通过Service访问，而是直接访问Pod，则可以使用此网段。
attachedVPC	Yes	String	Network所在VPC的ID。
networkType	Yes	String	Network网络类型，VPC网络类型取值： underlay_neutron 。
networkID	Yes	String	Network对应VPC子网的网络ID。
subnetID	Yes	String	Network对应VPC子网的子网ID。

参数	是否必选	参数类型	描述
availableZone	Yes	String	Network对应VPC子网所在可用区。当前支持“华北-北京四”、“华东-上海一”、“华东-上海二”和“华南-广州”区域，该值可设置为、“cn-north-4a”、“cn-east-3a”、“cn-east-2d”或“cn-south-1f”。

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "kind": "Network",
  "metadata": {
    "annotations": {
      "network.alpha.kubernetes.io/default-security-group": "security-group-id",
      "network.alpha.kubernetes.io/domain-id": "domain-id",
      "network.alpha.kubernetes.io/project-id": "project-id"
    },
    "name": "test-network",
    "namespace": "test-ns"
  },
  "spec": {
    "attachedVPC": "vpc-id",
    "availableZone": "cn-north-4a",
    "cidr": "192.168.0.0/24",
    "networkID": "network-id",
    "networkType": "underlay_neutron",
    "subnetID": "subnet-id"
  }
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "kind": "Network",
  "metadata": {
    "annotations": {
      "network.alpha.kubernetes.io/default-security-group": "security-group-id",
      "network.alpha.kubernetes.io/domain-id": "domain-id",
      "network.alpha.kubernetes.io/project-id": "project-id",
      "network.alpha.kubernetes.io/type": "underlay_neutron"
    },
    "creationTimestamp": "2018-08-21T02:35:59Z",
    "name": "test-network",
    "namespace": "test-ns",
    "resourceVersion": "2025736",
    "selfLink": "/apis/networking.cci.io/v1beta1/namespaces/test-ns/networks/test-network",
    "uid": "f03452ac-a4ea-11e8-8500-c81fbc371a17"
  },
  "spec": {
```

```
"attachedVPC" : "vpc-id",
"availableZone" : "cn-north-4a",
"cidr" : "192.168.0.0/24",
"networkID" : "network-id",
"networkType" : "underlay_neutron",
"subnetID" : "subnet-id"
},
"status" : {
  "state" : "Initializing"
}
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.3.4 删除 Network

功能介绍

删除一个指定Network对象。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks/{name}

表 4-248 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Network
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-249 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-250 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 4-251 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 4-252 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 4-253 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-254 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 4-255 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-256 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-257 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-258 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-259 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

状态码： 202

表 4-260 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-261 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-262 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-263 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-264 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-265 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-266 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

请求示例

删除Network。

```
{
  "Kind": "DeleteOptions",
  "apiVersion": "v1",
  "gracePeriodSeconds": 0
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "kind": "Network",
  "metadata": {
```

```
"annotations": {
  "network.alpha.kubernetes.io/default-security-group": "19c5d024-aed5-4856-b958-c0f65ce70855",
  "network.alpha.kubernetes.io/domain-id": "aad43c0b14c4cafbcfff483d075987",
  "network.alpha.kubernetes.io/project-id": "51bf52609f2a49c68bfda3398817b376"
},
"creationTimestamp": "2018-09-03T11:21:00Z",
"name": "namespace-test-dc1-default-network",
"namespace": "namespace-test",
"resourceVersion": "5016899",
"selfLink": "/apis/networking.cci.io/v1beta1/namespaces/namespace-test/networks/namespace-test-dc1-default-network",
"uid": "6fb85414-af6b-11e8-b6ef-f898ef6c78b4"
},
"spec": {
  "attachedVPC": "0d4080e5-546a-46c4-86fe-f3e26d685177",
  "availableZone": "cn-north-4a",
  "cidr": "192.168.244.0/23",
  "networkID": "0022e356-f730-4226-802e-9cdaa6e7da17",
  "networkType": "underlay_neutron",
  "subnetID": "1ffd839d-e534-4fa8-a59d-42356335bf74"
},
"status": {
  "state": "Active"
}
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.3.5 查询 Network

功能介绍

查询指定Network对象。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks/{name}

表 4-267 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Network
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-268 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-269 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-270 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-271 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-272 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-273 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-274 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-275 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-276 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "kind": "Network",
  "metadata": {
    "annotations": {
      "network.alpha.kubernetes.io/default-security-group": "19c5d024-aed5-4856-b958-c0f65ce70855",
      "network.alpha.kubernetes.io/domain-id": "aad43c0b14c4afbccfff483d075987",
      "network.alpha.kubernetes.io/project-id": "51bf52609f2a49c68bfda3398817b376"
    }
  },
  "creationTimestamp": "2018-09-03T11:21:00Z",
```



```
"name" : "namespace-test-dc1-default-network",
"namespace" : "namespace-test",
"resourceVersion" : "5016899",
"selfLink" : "/apis/networking.cci.io/v1beta1/namespaces/namespace-test/networks/namespace-test-dc1-
default-network",
"uid" : "6fb85414-af6b-11e8-b6ef-f898ef6c78b4"
},
"spec" : {
"attachedVPC" : "0d4080e5-546a-46c4-86fe-f3e26d685177",
"availableZone" : "cn-north-4a",
"cidr" : "192.168.244.0/23",
"networkID" : "0022e356-f730-4226-802e-9cdaa6e7da17",
"networkType" : "underlay_neutron",
"subnetID" : "1ffd839d-e534-4fa8-a59d-42356335bf74"
},
"status" : {
"state" : "Active"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

4.3.6 查询 Network 状态

功能介绍

查询一个指定Network对象的状态。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks/{name}/status

表 4-277 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Network
namespace	是	String	object name and auth scope, such as for teams and projects

表 4-278 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 4-279 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 4-280 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.cci.v1beta1.NetworkSpec object	Spec defines the attributes on a network
status	io.k8s.api.cci.v1beta1.NetworkStatus object	Status describes the network status

表 4-281 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-282 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 4-283 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 4-284 io.k8s.api.cci.v1beta1.NetworkSpec

参数	参数类型	描述
attachedVPC	String	ID of the VPC to attach
availableZone	String	available zone
cidr	String	The CIDR of the network

参数	参数类型	描述
networkID	String	network ID
networkType	String	network type like 'overlay_l2'
subnetID	String	Subnet ID
subnets	Array of io.k8s.api.cci.v1beta1.NetworkSubnet objects	Subnets

表 4-285 io.k8s.api.cci.v1beta1.NetworkSubnet

参数	参数类型	描述
cidr	String	The CIDR of the network
networkID	String	NetworkID describes the subnet network id
subnetID	String	SubnetID describes the subnet id

表 4-286 io.k8s.api.cci.v1beta1.NetworkStatus

参数	参数类型	描述
message	String	Message describes why network is in current state
state	String	State describes the network state

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "networking.cci.io/v1beta1",
  "kind": "Network",
  "metadata": {
    "annotations": {
      "network.alpha.kubernetes.io/default-security-group": "security-group-id",
      "network.alpha.kubernetes.io/domain-id": "domain-id",
      "network.alpha.kubernetes.io/project-id": "project-id",
      "network.alpha.kubernetes.io/type": "underlay_neutron"
    }
  },
}
```

```
"creationTimestamp" : "2018-08-21T02:35:59Z",
"name" : "test-network",
"namespace" : "test-ns",
"resourceVersion" : "2025736",
"selfLink" : "/apis/networking.cci.io/v1beta1/namespaces/test-ns/networks/test-network",
"uid" : "f03452ac-a4ea-11e8-8500-c81fbe371a17"
},
"spec" : {
  "attachedVPC" : "vpc-id",
  "availableZone" : "cn-north-4a",
  "cidr" : "192.168.0.0/24",
  "networkID" : "network-id",
  "networkType" : "underlay_neutron",
  "subnetID" : "subnet-id"
},
"status" : {
  "state" : "Active"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5 Kubernetes API

5.1 ConfigMap

5.1.1 删除指定 namespace 下的 ConfigMaps

功能介绍

删除Namespace下所有ConfigMap。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/configmaps

表 5-1 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-5 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-6 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-7 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-8 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-9 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "data": {
      "upstreamNameservers": "[\"*.*.*.*\"]"
    },
    "metadata": {
      "creationTimestamp": "2018-09-03T11:20:54Z",
      "labels": {
        "addonmanager.kubernetes.io/mode": "EnsureExists",
        "app": "kube-dns"
      },
      "name": "kube-dns",
      "namespace": "namespace-test",
      "resourceVersion": "5016780",
      "selfLink": "/api/v1/namespaces/namespace-test/configmaps/kube-dns",
      "uid": "6c48d677-af6b-11e8-b6ef-f898ef6c78b4"
    }
  } ],
  "kind": "ConfigMapList",
  "metadata": {
    "resourceVersion": "5174188",
    "selfLink": "/api/v1/namespaces/namespace-test/configmaps"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.1.2 查询指定 namespace 下的 ConfigMaps

功能介绍

查询Namespace下所有ConfigMap的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/configmaps

表 5-10 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-11 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-12 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-13 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.ConfigMap objects	Items is the list of ConfigMaps.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-14 io.k8s.api.core.v1.ConfigMap

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '!'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '!'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-15 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-16 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-17 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-18 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "items": [{
    "data": {
      "property_1": "test"
    },
    "metadata": {
      "creationTimestamp": "2017-12-13T03:15:57Z",
      "name": "test-12130306",
      "namespace": "ns-12130306-s",
      "resourceVersion": "419081",
      "selfLink": "/api/v1/namespaces/ns-12130306-s/configmaps/test-12130306",
      "uid": "efd6d9e0-dfb3-11e7-9c19-fa163e2d897b"
    }
  ]},
  "kind": "ConfigMapList",
  "metadata": {
    "resourceVersion": "419140",
    "selfLink": "/api/v1/namespaces/ns-12130306-s/configmaps"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable

状态码	描述
504	ServerTimeout

5.1.3 创建 ConfigMap

功能介绍

创建ConfigMap。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/configmaps

表 5-19 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-20 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-21 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-22 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	否	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '!'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.

参数	是否必选	参数类型	描述
data	否	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '!'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	否	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-23 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-24 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-25 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

响应参数

状态码： 200

表 5-26 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '!'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '!'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-27 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-28 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-29 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

状态码： 201

表 5-30 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '.'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-31 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-32 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-33 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

状态码： 202

表 5-34 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '.'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-35 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-36 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-37 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

创建一个名称为“configmap-test”的configmap，配置数据键值。

```
{
  "apiVersion": "v1",
  "data": {
    "property_1": "test"
  },
  "kind": "ConfigMap",
  "metadata": {
    "name": "configmap-test"
  }
}
```

```
}  
}
```

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "v1",  
  "data": {  
    "property_1": "test"  
  },  
  "kind": "ConfigMap",  
  "metadata": {  
    "creationTimestamp": "2018-09-04T03:11:29Z",  
    "name": "configmap-test",  
    "namespace": "namespace-test",  
    "resourceVersion": "5170290",  
    "selfLink": "/api/v1/namespaces/namespace-test/configmaps/configmap-test",  
    "uid": "379519a3-aff0-11e8-8f17-c81fbe371a17"  
  }  
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.1.4 删除 ConfigMap

功能介绍

删除ConfigMap。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/configmaps/{name}

表 5-38 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ConfigMap
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-39 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-40 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-41 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-42 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-43 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-44 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-45 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-46 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-47 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-48 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none"> for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-49 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-50 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "code": 200,
  "details": {
    "kind": "configmaps",

```

```
"name" : "configmap-test",  
"uid" : "379519a3-aff0-11e8-8f17-c81fbe371a17"  
},  
"kind" : "Status",  
"metadata" : { },  
"status" : "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.1.5 查询 ConfigMap

功能介绍

查询ConfigMap详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/configmaps/{name}

表 5-51 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ConfigMap
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-52 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-53 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-54 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '.'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-55 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-56 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-57 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

无

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "v1",
```

```
"data" : {  
  "property_1" : "new"  
},  
"kind" : "ConfigMap",  
"metadata" : {  
  "creationTimestamp" : "2017-12-13T03:15:57Z",  
  "name" : "test-12130306",  
  "namespace" : "ns-12130306-s",  
  "resourceVersion" : "419141",  
  "selfLink" : "/api/v1/namespaces/ns-12130306-s/configmaps/test-12130306",  
  "uid" : "efd6d9e0-dfb3-11e7-9c19-fa163e2d897b"  
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.1.6 更新 ConfigMap

功能介绍

更新ConfigMap。

The following fields can be updated:

- metadata.labels
- metadata.annotations
- data

调用方法

请参见[如何调用API](#)。

URI

PATCH /api/v1/namespaces/{namespace}/configmaps/{name}

表 5-58 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ConfigMap
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-59 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

参数	是否必选	参数类型	描述
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-60 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-61 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-62 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '.'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-63 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-64 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-65 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

更新ConfigMap，添加一个property_2的键值对。

```
{
  "data": {
    "property_2": "test"
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "data": {
    "property_1": "test",
    "property_2": "test"
  },
  "kind": "ConfigMap",
  "metadata": {
    "creationTimestamp": "2018-09-04T03:11:29Z",
    "name": "configmap-test",
    "namespace": "namespace-test",
    "resourceVersion": "5171481",
    "selfLink": "/api/v1/namespaces/namespace-test/configmaps/configmap-test",
    "uid": "379519a3-aff0-11e8-8f17-c81fbe371a17"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.1.7 替换 ConfigMap

功能介绍

替换ConfigMap。

The following fields can be updated:

- metadata.labels
- metadata.annotations
- data

调用方法

请参见[如何调用API](#)。

URI

PUT /api/v1/namespaces/{namespace}/configmaps/{name}

表 5-66 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ConfigMap
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-67 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-68 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-69 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	否	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '!'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.

参数	是否必选	参数类型	描述
data	否	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '!'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	否	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-70 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-71 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-72 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

响应参数

状态码： 200

表 5-73 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '!'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '!'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-74 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-75 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-76 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

状态码： 201

表 5-77 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
binaryData	Map<String,String>	BinaryData contains the binary data. Each key must consist of alphanumeric characters, '-', '_' or '.'. BinaryData can contain byte sequences that are not in the UTF-8 range. The keys stored in BinaryData must not overlap with the ones in the Data field, this is enforced during validation process. Using this field will require 1.10+ apiserver and kubelet.
data	Map<String,String>	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. Values with non-UTF-8 byte sequences must use the BinaryData field. The keys stored in Data must not overlap with the keys in the BinaryData field, this is enforced during validation process.
immutable	Boolean	Immutable, if set to true, ensures that data stored in the ConfigMap cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-78 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-79 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-80 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

将已创建ConfigMap中的data数据替换为 "property_1" : "test2" 。

```
{
  "apiVersion": "v1",
  "data": {
    "property_1": "test2"
  },
  "kind": "ConfigMap",
  "metadata": {
    "name": "configmap-test"
```



```
}  
}
```

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "v1",  
  "data": {  
    "property_1": "test2"  
  },  
  "kind": "ConfigMap",  
  "metadata": {  
    "creationTimestamp": "2018-09-04T03:11:29Z",  
    "name": "configmap-test",  
    "namespace": "namespace-test",  
    "resourceVersion": "5172849",  
    "selfLink": "/api/v1/namespaces/namespace-test/configmaps/configmap-test",  
    "uid": "379519a3-aff0-11e8-8f17-c81fbe371a17"  
  }  
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2 Pod

5.2.1 删除指定 namespace 下的 Pods

功能介绍

删除Namespace下所有Pod。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/pods

表 5-81 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-82 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-83 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-84 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-85 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-86 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-87 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-88 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-89 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [],
  "kind": "PodList",
  "metadata": {
    "resourceVersion": "5035636",
    "selfLink": "/api/v1/namespaces/namespace-test/pods"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.2 查询指定 namespace 下的 Pods

功能介绍

查询指定namespace下的Pods的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/pods

表 5-90 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-91 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-92 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-93 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Pod objects	List of pods. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-94 io.k8s.api.core.v1.Pod

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-95 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	参数类型	描述
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.

参数	参数类型	描述
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-96 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-97 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-98 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-99 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-100 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-101 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-102 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-103 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-104 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-105 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-106 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-107 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-108 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-109 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-110 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-111 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-112 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-113 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <i>metadata.name</i> , <i>metadata.namespace</i> , <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , <i>spec.nodeName</i> , <i>spec.serviceAccountName</i> , <i>status.hostIP</i> , <i>status.podIP</i> , <i>status.podIPs</i> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-114 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-115 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-116 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-117 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-118 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-119 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-120 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-121 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-122 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-123 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-124 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-125 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-126 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-127 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-128 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-129 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-130 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-131 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-132 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-133 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-134 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-135 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-136 io.k8s.api.core.v1.SecompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-137 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-138 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-139 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-140 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-141 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-142 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-143 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-144 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-145 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-146 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-147 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-148 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-149 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-150 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-151 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-152 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-153 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-154 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-155 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-156 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-157 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-158 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-159 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-160 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-161 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-162 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-163 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-164 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-165 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-166 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-167 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-168 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-169 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-170 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-171 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-172 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-173 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-174 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-175 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-176 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-177 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-178 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-179 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-180 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-181 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-182 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-183 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-184 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-185 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-186 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-187 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-188 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-189 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-190 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-191 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-192 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-193 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-194 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-195 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-196 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-197 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

表 5-198 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "annotations": {
        "cri.cci.io/container-type": "secure-container",
        "kubernetes.io/availablezone": "dc1",
        "network.alpha.kubernetes.io/network": "[{\"name\":\"namespace-test-dc1-default-network\", \"interface\":\"eth0\", \"network_plane\":\"default\"}]"
      },
      "creationTimestamp": "2018-09-03T12:26:12Z",
      "labels": {
        "name": "pod-test"
      },
      "name": "pod-test",
      "namespace": "namespace-test",
      "resourceVersion": "5030610",
      "selfLink": "/api/v1/namespaces/namespace-test/pods/pod-test",
      "uid": "8b985a27-af74-11e8-9d5d-c88d83be759f"
    },
    "spec": {
      "containers": [ {
        "image": "redis",
        "imagePullPolicy": "Always",
        "name": "test",
        "resources": {
          "limits": {
            "cpu": "500m",
            "memory": "1Gi"
          },
          "requests": {
            "cpu": "500m",
            "memory": "1Gi"
          }
        },
        "terminationMessagePath": "/dev/termination-log",
        "terminationMessagePolicy": "File"
      } ],
      "dnsPolicy": "ClusterFirst",
      "imagePullSecrets": [ {
        "name": "imagepull-secret"
      } ],
      "nodeName": "c0dd6256-195a-e811-90a2-10c17294fcbc",
      "restartPolicy": "Always",
      "schedulerName": "default-scheduler",
      "securityContext": { },
      "tolerations": [ {
        "effect": "NoExecute",
        "key": "node.kubernetes.io/not-ready",
        "operator": "Exists",
        "tolerationSeconds": 300
      }, {
        "effect": "NoExecute",
        "key": "node.kubernetes.io/unreachable",
        "operator": "Exists",
        "tolerationSeconds": 300
      } ]
    }
  } ],
}
```

```
"status": {
  "conditions": [ {
    "lastProbeTime": null,
    "lastTransitionTime": "2018-09-03T12:26:12Z",
    "status": "True",
    "type": "Initialized"
  }, {
    "lastProbeTime": null,
    "lastTransitionTime": "2018-09-03T12:26:16Z",
    "status": "True",
    "type": "Ready"
  }, {
    "lastProbeTime": null,
    "lastTransitionTime": "2018-09-03T12:26:12Z",
    "status": "True",
    "type": "PodScheduled"
  } ],
  "containerStatuses": [ {
    "containerID": "docker://aee55d8dedb8371f96aa5d5116f69a53bf1cb23afe1802567c24081514d3b048",
    "image": "redis",
    "imageID": "docker-pullable://
redis@sha256:3ab7046bd035a47aa06963d8240651d00b57e82dab07ba374ad01f84dfa1230c",
    "lastState": { },
    "name": "test",
    "ready": true,
    "restartCount": 0,
    "state": {
      "running": {
        "startedAt": "2018-09-03T12:26:16Z"
      }
    }
  } ],
  "phase": "Running",
  "podIP": "192.168.245.185",
  "qosClass": "Guaranteed",
  "startTime": "2018-09-03T12:26:12Z"
} ],
"kind": "PodList",
"metadata": {
  "resourceVersion": "5032373",
  "selfLink": "/api/v1/namespaces/namespace-test/pods"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.3 创建 Pod

功能介绍

创建一个Pod。

说明

说明：Pod的生命周期是短暂的，Pod是用后即焚的实体。在实际使用中，请谨慎单独创建Pod，请使用Deployment、StatefulSet和Job这些控制器创建应用，从而保障应用高可用。

云容器实例中Pod规格有限制，具体的限制请参见[约束限制](#)页面的“Pod规格”部分。

当前支持使用Nvidia GPU的驱动版本为418.126和460.106，您应用程序中使用的CUDA需满足如下表所示的配套关系。CUDA与驱动的配套关系来源于Nvidia官网，详细信息请参见[CUDA Compatibility](#)。

表：NVIDIA GPU驱动与CUDA配套关系

NVIDIA GPU驱动版本	CUDA Toolkit版本
418.126	CUDA 10.1 (10.1.105)及以下
460.106	CUDA 11.2.2 Update 2 及以下

如果选择GPU加速型Pod，您需要设置Pod的**metadata.annotations**中添加**cri.cci.io/gpu-driver**字段，指定使用哪个版本显卡驱动，取值如下：

- gpu-418.126
- gpu-460.106

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/pods

表 5-199 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-200 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-201 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-202 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-203 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.

参数	是否必选	参数类型	描述
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.

参数	是否必选	参数类型	描述
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	是否必选	参数类型	描述
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.

参数	是否必选	参数类型	描述
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	是否必选	参数类型	描述
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceName	否	String	DeprecatedServiceAccount is a depreciated alias for ServiceAccountName. Deprecated: Use serviceName instead.

参数	是否必选	参数类型	描述
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.

参数	是否必选	参数类型	描述
subdomain	否	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	是否必选	参数类型	描述
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-204 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-205 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-206 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	是否必选	参数类型	描述
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-207 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-208 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-209 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	是否必选	参数类型	描述
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-210 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-211 `io.k8s.api.core.v1.PodAntiAffinity`

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-212 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-213 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"

参数	是否必选	参数类型	描述
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-214 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-215 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.

参数	是否必选	参数类型	描述
value	否	String	value is the value of the option

表 5-216 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	是否必选	参数类型	描述
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	是否必选	参数类型	描述
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-217 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-218 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-219 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-220 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-221 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-222 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-223 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-224 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-225 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-226 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-227 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-228 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-229 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-230 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-231 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-232 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-233 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-234 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-235 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-236 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-237 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-238 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain '!'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-239 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-240 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-241 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-242 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-243 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-244 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-245 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-246 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	否	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	否	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-247 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-248 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-249 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-250 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-251 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-252 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-253 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-254 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-255 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-256 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-257 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-258 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-259 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-260 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-261 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-262 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-263 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-264 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-265 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-266 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-267 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-268 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-269 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-270 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-271 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-272 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-273 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-274 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-275 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-276 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <code>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br]</code> will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.

参数	是否必选	参数类型	描述
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-277 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) (Note that [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) may be empty, from the "" case in [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= 0 1 ... 9 [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br). ./[topic/body/section/</p>

参数	是否必选	参数类型	描述
			<p>table/tgroup/tbody/row/entry/p/br {""} (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "+" "-" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2^63-1</p>

参数	是否必选	参数类型	描述
			<p>in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-278 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-279 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-280 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-281 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-282 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-283 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-284 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-285 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-286 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-287 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-288 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-289 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-290 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-291 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-292 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-293 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-294 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-295 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-296 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-297 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmdk

表 5-298 io.k8s.api.core.v1.PodStatus

参数	是否必选	参数类型	描述
conditions	否	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	否	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	否	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	否	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	否	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	否	String	A human readable message indicating details about why the pod is in this condition.

参数	是否必选	参数类型	描述
nominatedNodeName	否	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.

参数	是否必选	参数类型	描述
phase	否	String	<p>The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values:</p> <p>Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod.</p> <p>More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase</p>
podIP	否	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.

参数	是否必选	参数类型	描述
podIPs	否	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.
qosClass	否	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	否	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	否	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-299 io.k8s.api.core.v1.PodCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, one-word, CamelCase reason for the condition's last transition.

参数	是否必选	参数类型	描述
status	是	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	是	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-300 io.k8s.api.core.v1.ContainerStatus

参数	是否必选	参数类型	描述
containerID	否	String	Container's ID in the format 'docker://<container_id>'.
image	是	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images
imageID	是	String	ImageID of the container's image.
lastState	否	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	是	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	是	Boolean	Specifies whether the container has passed its readiness probe.

参数	是否必选	参数类型	描述
restartCount	是	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	否	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	否	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-301 io.k8s.api.core.v1.ContainerState

参数	是否必选	参数类型	描述
running	否	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	否	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	否	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-302 io.k8s.api.core.v1.ContainerStateRunning

参数	是否必选	参数类型	描述
startedAt	否	String	Time at which the container was last (re-)started

表 5-303 io.k8s.api.core.v1.ContainerStateTerminated

参数	是否必选	参数类型	描述
containerID	否	String	Container's ID in the format 'docker://<container_id>'
exitCode	是	Integer	Exit status from the last termination of the container
finishedAt	否	String	Time at which the container last terminated
message	否	String	Message regarding the last termination of the container
reason	否	String	(brief) reason from the last termination of the container
signal	否	Integer	Signal from the last termination of the container
startedAt	否	String	Time at which previous execution of the container started

表 5-304 io.k8s.api.core.v1.ContainerStateWaiting

参数	是否必选	参数类型	描述
message	否	String	Message regarding why the container is not yet running.
reason	否	String	(brief) reason the container is not yet running.

表 5-305 io.k8s.api.core.v1.PodIP

参数	是否必选	参数类型	描述
ip	否	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

响应参数

状态码： 200

表 5-306 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-307 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.

参数	参数类型	描述
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-308 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-309 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-310 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-311 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-312 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-313 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-314 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-315 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-316 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-317 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-318 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-319 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-320 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-321 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-322 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-323 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-324 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-325 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-326 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-327 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-328 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-329 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-330 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-331 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-332 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-333 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-334 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-335 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-336 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-337 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-338 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-339 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-340 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-341 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-342 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-343 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-344 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-345 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-346 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-347 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-348 io.k8s.api.core.v1.SecompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-349 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-350 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-351 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-352 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-353 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-354 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-355 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-356 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-357 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-358 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-359 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-360 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-361 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-362 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-363 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-364 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-365 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-366 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-367 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-368 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-369 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-370 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-371 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-372 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-373 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-374 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-375 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-376 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-377 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-378 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-379 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-380 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-381 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre> [/<topic ""="" "-"="" (br)="" (br).="" (br).)="" (br)[="" (international="" (note="" .="" ...="" 1="" 9="" ::="Ki" <a="" <topic="" [="" be="" body="" br="" case="" ei="" empty,="" entry="" from="" gi="" href="http://physics.nist.gov/cuu/Units/binary.html" in="" may="" mi="" of="" p="" pi="" row="" section="" see:="" system="" table="" tbody="" tgroup="" that="" the="" ti="" units;="" {""}="" ="">http://physics.nist.gov/cuu/Units/binary.html) [/<topic <="" body="" pre="" section="" table="" tgroup=""> </topic></topic></pre>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-382 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-383 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-384 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-385 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-386 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-387 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-388 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-389 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-390 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-391 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-392 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-393 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-394 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-395 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-396 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-397 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-398 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-399 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-400 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-401 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-402 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-403 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-404 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-405 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-406 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-407 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-408 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-409 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

状态码: 201

表 5-410 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-411 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-412 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-413 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-414 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-415 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-416 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-417 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-418 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-419 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-420 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-421 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-422 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-423 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-424 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-425 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-426 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-427 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-428 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-429 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-430 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-431 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-432 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-433 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-434 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-435 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-436 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-437 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-438 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-439 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-440 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-441 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-442 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-443 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-444 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-445 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-446 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-447 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-448 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-449 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-450 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-451 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-452 io.k8s.api.core.v1.SecompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-453 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-454 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-455 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-456 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-457 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-458 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-459 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-460 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-461 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-462 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-463 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-464 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-465 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-466 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-467 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-468 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-469 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-470 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-471 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-472 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-473 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-474 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/

表 5-475 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-476 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-477 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-478 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-479 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-480 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-481 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-482 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-483 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-484 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-485 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) (Note that [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) may be empty, from the "" case in [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= 0 1 ... 9 [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br). .[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "+" "-" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [/topic/body/section/table/tgroup/</p>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-486 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-487 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-488 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-489 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-490 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-491 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-492 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-493 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-494 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-495 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-496 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-497 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-498 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-499 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-500 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-501 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-502 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-503 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-504 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-505 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-506 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-507 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-508 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-509 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-510 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-511 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-512 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-513 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

状态码: 202

表 5-514 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-515 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-516 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-517 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-518 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-519 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-520 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-521 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-522 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-523 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-524 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-525 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-526 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-527 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-528 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-529 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-530 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-531 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-532 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-533 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-534 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-535 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-536 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-537 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-538 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-539 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-540 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-541 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-542 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-543 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-544 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-545 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-546 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-547 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-548 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-549 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-550 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-551 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-552 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-553 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-554 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-555 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-556 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-557 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-558 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-559 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-560 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-561 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-562 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-563 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-564 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-565 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-566 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-567 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-568 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-569 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-570 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-571 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-572 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-573 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-574 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-575 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-576 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-577 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-578 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-579 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-580 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-581 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-582 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-583 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-584 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-585 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-586 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-587 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-588 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-589 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) (Note that [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) may be empty, from the "" case in [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= 0 1 ... 9 [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br). .[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "+" "-" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [/topic/body/section/table/tgroup/</p>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-590 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-591 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-592 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-593 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-594 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-595 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-596 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-597 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-598 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-599 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-600 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-601 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-602 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-603 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-604 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-605 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-606 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-607 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-608 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-609 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-610 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-611 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-612 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-613 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-614 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-615 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-616 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-617 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

请求示例

- 创建普通Pod，使用redis镜像，Pod资源占用CPU为0.5，内存为1024Mi。

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "labels": {
      "name": "pod-test"
    },
    "name": "pod-test"
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "0.5",
          "memory": "1024Mi"
        },
        "requests": {
          "cpu": "0.5",
          "memory": "1024Mi"
        }
      }
    }
  ],
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "priority": 0,
  "restartPolicy": "Always"
}
```

- 创建GPU型Pod，使用gpu-418.126版本显示驱动，使用redis镜像，Pod资源占用CPU为4，内存为32Gi。

📖 说明

说明：“华北-北京四”区域，仅支持NVIDIA TeslaV100 32G显卡。“华东-上海一”可支持NVIDIA TeslaV100 32G、NVIDIA TeslaV100 16G显卡。

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "annotations": {
      "cri.cci.io/gpu-driver": "gpu-418.126"
    },
    "labels": {
      "name": "pod-test"
    },
    "name": "pod-test"
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "4",
          "memory": "32Gi",
          "nvidia.com/gpu-tesla-v100-16GB": "1"
        },
        "requests": {
          "cpu": "4",
          "memory": "32Gi",

```

```
    "nvidia.com/gpu-tesla-v100-16GB" : "1"
  }
}
}],
"imagePullSecrets" : [ {
  "name" : "imagepull-secret"
}],
"priority" : 0,
"restartPolicy" : "Always"
}
}
```

- 通过LVM方式创建本地卷的Pod，容器内卷的挂载路径为“/tmp/log”。

```
{
  "apiVersion" : "v1",
  "kind" : "Pod",
  "metadata" : {
    "labels" : {
      "app" : "localvolume"
    },
  },
  "name" : "localvolume",
  "namespace" : "cci-namespace-12192721"
},
"spec" : {
  "containers" : [ {
    "command" : [ "/bin/sh", "-c", "sleep 10000" ],
    "image" : "100.79.1.215:20202/paas_cci_test/redis:v1",
    "imagePullPolicy" : "Always",
    "lifecycle" : { },
    "name" : "container-0",
    "resources" : {
      "limits" : {
        "cpu" : "500m",
        "memory" : "1024Mi"
      },
      "requests" : {
        "cpu" : "500m",
        "memory" : "1024Mi"
      }
    },
  },
  "terminationMessagePath" : "/dev/termination-log",
  "terminationMessagePolicy" : "File",
  "volumeMounts" : [ {
    "mountPath" : "/tmp/log",
    "name" : "innerevs123"
  } ]
}],
"dnsPolicy" : "Default",
"imagePullSecrets" : [ {
  "name" : "imagepull-secret"
}],
"nodeSelector" : {
  "node.cci.io/allowed-on-poc-dedicated-node" : "sina"
},
"restartPolicy" : "Always",
"schedulerName" : "default-scheduler",
"tolerations" : [ {
  "effect" : "NoSchedule",
  "key" : "node.cci.io/allowed-on-poc-dedicated-node",
  "operator" : "Equal",
  "value" : "sina"
}],
"volumes" : [ {
  "emptyDir" : {
    "sizeLimit" : "100Gi"
  },
  "name" : "innerevs123"
} ]
}
}
```

响应示例

状态码： 201

Created

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "annotations": {
      "cri.cci.io/container-type": "secure-container"
    },
    "creationTimestamp": "2022-08-30T03:14:15Z",
    "finalizers": [ "billing.kubernetes.io/podbilling-finalizer" ],
    "labels": {
      "name": "pod-test",
      "tenant.kubernetes.io/domain-id": "xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx"
    },
    "name": "pod-test",
    "namespace": "test",
    "resourceVersion": "33742743",
    "selfLink": "/api/v1/namespaces/test/pods/pod-test",
    "uid": "c0f9820e-188c-4612-be69-b9c65c8bdb5a"
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      }
    }
  ],
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
},
  "dnsPolicy": "Default",
  "enableServiceLinks": false,
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "nodeSelector": {
    "node.cci.io/default-cpu-choice": "true",
    "node.cci.io/flavor": "general-computing"
  },
  "priority": 0,
  "restartPolicy": "Always",
  "runtimeClassName": "kata",
  "schedulerName": "volcano",
  "securityContext": { },
  "terminationGracePeriodSeconds": 30,
  "tolerations": [ {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/not-ready",
    "operator": "Exists",
    "tolerationSeconds": 300
  }, {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/unreachable",
    "operator": "Exists",
    "tolerationSeconds": 300
  }, {
```

```
"effect" : "NoSchedule",
"key" : "node.cci.io/occupied",
"operator" : "Equal",
"value" : "default"
}, {
"effect" : "NoSchedule",
"key" : "node.cci.io/allowed-on-shared-node",
"operator" : "Exists"
}, {
"effect" : "NoSchedule",
"key" : "node.kubernetes.io/memory-pressure",
"operator" : "Exists"
}
}
},
"status" : {
"phase" : "Pending",
"qosClass" : "Guaranteed"
}
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.4 删除 Pod

功能介绍

删除Pod，删除前持续10s。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/pods/{name}

表 5-618 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Pod
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-619 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-620 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-621 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-622 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-623 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-624 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	参数类型	描述
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-625 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-626 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-627 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-628 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-629 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-630 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-631 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-632 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-633 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-634 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-635 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-636 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-637 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-638 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-639 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-640 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-641 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-642 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-643 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-644 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-645 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-646 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-647 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-648 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-649 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-650 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-651 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-652 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-653 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-654 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-655 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-656 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-657 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-658 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-659 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-660 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-661 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-662 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-663 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-664 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-665 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-666 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-667 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-668 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-669 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-670 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-671 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-672 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-673 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-674 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-675 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-676 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-677 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-678 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-679 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-680 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-681 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-682 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-683 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-684 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-685 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-686 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-687 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-688 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-689 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-690 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-691 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-692 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-693 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-694 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-695 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-696 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-697 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-698 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-699 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-700 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-701 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-702 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-703 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-704 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-705 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-706 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-707 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-708 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-709 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-710 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-711 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-712 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-713 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-714 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-715 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-716 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-717 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-718 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-719 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-720 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-721 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-722 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-723 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-724 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-725 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-726 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

状态码: 202

表 5-727 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-728 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	参数类型	描述
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-729 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-730 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-731 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-732 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-733 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-734 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-735 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-736 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-737 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-738 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-739 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-740 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-741 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-742 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-743 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-744 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-745 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-746 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-747 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-748 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-749 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-750 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-751 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-752 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-753 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-754 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-755 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-756 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-757 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-758 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-759 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-760 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-761 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-762 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-763 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-764 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-765 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-766 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-767 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-768 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-769 io.k8s.api.core.v1.SecompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-770 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-771 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-772 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-773 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-774 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	<p>FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.</p>
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	<p>Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running</p>
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	<p>GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</p>

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-775 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-776 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-777 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-778 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-779 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-780 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-781 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-782 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-783 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-784 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-785 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-786 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-787 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-788 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-789 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-790 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-791 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-792 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-793 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-794 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-795 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-796 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-797 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-798 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-799 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-800 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-801 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-802 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-803 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-804 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-805 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-806 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-807 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-808 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-809 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-810 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-811 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-812 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-813 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-814 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-815 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-816 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-817 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-818 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-819 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-820 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-821 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-822 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-823 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-824 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-825 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-826 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-827 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-828 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-829 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-830 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

请求示例

```
{  
  "apiVersion": "v1",
```

```
"gracePeriodSeconds": 10,  
"kind": "DeleteOptions"  
}
```

响应示例

状态码： 200

OK

```
{  
  "apiVersion": "v1",  
  "items": [ {  
    "metadata": {  
      "annotations": {  
        "cri.cci.io/container-type": "secure-container",  
        "kubernetes.io/availablezone": "dc1",  
        "network.alpha.kubernetes.io/network": "[{\\"name\\":\\"namespace-test-dc1-default-network\\",\\"interface\\":\\"eth0\\",\\"network_plane\\":\\"default\\"}]"  
      },  
      "creationTimestamp": "2018-09-03T12:26:12Z",  
      "labels": {  
        "name": "pod-test"  
      },  
      "name": "pod-test",  
      "namespace": "namespace-test",  
      "resourceVersion": "5030610",  
      "selfLink": "/api/v1/namespaces/namespace-test/pods/pod-test",  
      "uid": "8b985a27-af74-11e8-9d5d-c88d83be759f"  
    },  
    "spec": {  
      "containers": [ {  
        "image": "redis",  
        "imagePullPolicy": "Always",  
        "name": "test",  
        "resources": {  
          "limits": {  
            "cpu": "500m",  
            "memory": "1Gi"  
          },  
          "requests": {  
            "cpu": "500m",  
            "memory": "1Gi"  
          }  
        },  
        "terminationMessagePath": "/dev/termination-log",  
        "terminationMessagePolicy": "File"  
      } ],  
      "dnsPolicy": "ClusterFirst",  
      "imagePullSecrets": [ {  
        "name": "imagepull-secret"  
      } ],  
      "nodeName": "c0dd6256-195a-e811-90a2-10c17294fcbc",  
      "restartPolicy": "Always",  
      "schedulerName": "default-scheduler",  
      "securityContext": { },  
      "tolerations": [ {  
        "effect": "NoExecute",  
        "key": "node.kubernetes.io/not-ready",  
        "operator": "Exists",  
        "tolerationSeconds": 300  
      }, {  
        "effect": "NoExecute",  
        "key": "node.kubernetes.io/unreachable",  
        "operator": "Exists",  
        "tolerationSeconds": 300  
      } ]  
    },  
    "status": {  
      "phase": "Running",  
      "podIP": "10.1.1.1",  
      "startTime": "2018-09-03T12:26:12Z"  
    }  
  } ]  
}
```

```
"conditions": [ {
  "lastProbeTime": null,
  "lastTransitionTime": "2018-09-03T12:26:12Z",
  "status": "True",
  "type": "Initialized"
}, {
  "lastProbeTime": null,
  "lastTransitionTime": "2018-09-03T12:26:16Z",
  "status": "True",
  "type": "Ready"
}, {
  "lastProbeTime": null,
  "lastTransitionTime": "2018-09-03T12:26:12Z",
  "status": "True",
  "type": "PodScheduled"
} ],
"containerStatuses": [ {
  "containerID": "docker://aee55d8dedb8371f96aa5d5116f69a53bf1cb23afe1802567c24081514d3b048",
  "image": "redis",
  "imageID": "docker-pullable://
redis@sha256:3ab7046bd035a47aa06963d8240651d00b57e82dab07ba374ad01f84dfa1230c",
  "lastState": { },
  "name": "test",
  "ready": true,
  "restartCount": 0,
  "state": {
    "running": {
      "startedAt": "2018-09-03T12:26:16Z"
    }
  }
} ],
"managementIP": "172.28.0.17",
"phase": "Running",
"podIP": "192.168.245.185",
"qosClass": "Guaranteed",
"startTime": "2018-09-03T12:26:12Z"
} ],
"kind": "PodList",
"metadata": {
  "resourceVersion": "5035636",
  "selfLink": "/api/v1/namespaces/namespace-test/pods"
}
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict

状态码	描述
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.5 查询 Pod

功能介绍

查询Pod的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/pods/{name}

表 5-831 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Pod
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-832 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.

参数	是否必选	参数类型	描述
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-833 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-834 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-835 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	参数类型	描述
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.

参数	参数类型	描述
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-836 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-837 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-838 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-839 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-840 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-841 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-842 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-843 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-844 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-845 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-846 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-847 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-848 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-849 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-850 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-851 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-852 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-853 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-854 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-855 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-856 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-857 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-858 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-859 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-860 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-861 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-862 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-863 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-864 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-865 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-866 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-867 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-868 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-869 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-870 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '..' and must not be a root directory.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-871 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-872 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-873 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-874 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-875 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-876 io.k8s.api.core.v1.SecompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-877 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-878 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-879 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-880 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-881 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-882 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-883 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-884 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-885 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-886 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-887 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-888 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-889 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-890 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-891 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-892 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-893 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-894 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-895 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-896 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-897 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-898 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-899 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-900 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-901 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-902 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-903 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-904 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-905 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-906 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-907 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-908 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-909 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-910 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-911 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-912 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-913 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-914 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-915 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-916 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-917 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-918 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-919 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-920 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-921 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-922 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-923 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-924 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-925 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-926 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-927 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-928 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-929 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-930 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-931 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-932 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-933 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-934 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-935 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-936 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-937 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "annotations": {
      "cri.cci.io/container-type": "secure-container",
      "kubernetes.io/availablezone": "dc1",
      "network.alpha.kubernetes.io/network": "[{\"name\": \"namespace-test-dc1-default-network\", \"interface\": \"eth0\", \"network_plane\": \"default\"}]"
    },
    "creationTimestamp": "2018-09-03T12:26:12Z",
    "labels": {
      "name": "pod-test"
    },
    "name": "pod-test",
    "namespace": "namespace-test",
    "resourceVersion": "5030610",
    "selfLink": "/api/v1/namespaces/namespace-test/pods/pod-test",
    "uid": "8b985a27-af74-11e8-9d5d-c88d83be759f"
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      }
    }
  ],
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
},
  "dnsPolicy": "ClusterFirst",
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "nodeName": "c0dd6256-195a-e811-90a2-10c17294fcbc",
  "restartPolicy": "Always",
  "schedulerName": "default-scheduler",
  "securityContext": { },
  "tolerations": [ {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/not-ready",
    "operator": "Exists",
    "tolerationSeconds": 300
  }, {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/unreachable",
    "operator": "Exists",
    "tolerationSeconds": 300
  } ]
},
  "status": {
    "conditions": [ {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-03T12:26:12Z",
      "status": "True",
    } ]
  }
}
```

```

    "type" : "Initialized"
  }, {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2018-09-03T12:26:16Z",
    "status" : "True",
    "type" : "Ready"
  }, {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2018-09-03T12:26:12Z",
    "status" : "True",
    "type" : "PodScheduled"
  } ],
  "containerStatuses" : [ {
    "containerID" : "docker://aee55d8dedb8371f96aa5d5116f69a53bf1cb23afe1802567c24081514d3b048",
    "image" : "redis",
    "imageID" : "docker-pullable://
redis@sha256:3ab7046bd035a47aa06963d8240651d00b57e82dab07ba374ad01f84dfa1230c",
    "lastState" : { },
    "name" : "test",
    "ready" : true,
    "restartCount" : 0,
    "state" : {
      "running" : {
        "startedAt" : "2018-09-03T12:26:16Z"
      }
    }
  } ],
  "phase" : "Running",
  "podIP" : "192.168.245.185",
  "qosClass" : "Guaranteed",
  "startTime" : "2018-09-03T12:26:12Z"
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.6 更新 Pod

功能介绍

更新Pod。

其中以下字段支持更新：

- metadata.generateName
- metadata.labels
- metadata.annotations
- spec.initContainers[*].image
- spec.containers[*].image
- spec.activeDeadlineSeconds
- spec.tolerations 其余部分不支持更新。

调用方法

请参见[如何调用API](#)。

URI

PATCH /api/v1/namespaces/{namespace}/pods/{name}

表 5-938 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Pod
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-939 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-940 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-941 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-942 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-943 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	参数类型	描述
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-944 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-945 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-946 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-947 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-948 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-949 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-950 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-951 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-952 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-953 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-954 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-955 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-956 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-957 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-958 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-959 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-960 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-961 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-962 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-963 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-964 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-965 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-966 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-967 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-968 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-969 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-970 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-971 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-972 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-973 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-974 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-975 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-976 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-977 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-978 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-979 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-980 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-981 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-982 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: 1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-983 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-984 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-985 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-986 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-987 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-988 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-989 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-990 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-991 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-992 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-993 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-994 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-995 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-996 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-997 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-998 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-999 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1000 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1001 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1002 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1003 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1004 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1005 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1006 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1007 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1008 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1009 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1010 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1011 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1012 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1013 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1014 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1015 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1016 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1017 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1018 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1019 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1020 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1021 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1022 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1023 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1024 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1025 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1026 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1027 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1028 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1029 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1030 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1031 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1032 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1033 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1034 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1035 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1036 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1037 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1038 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-1039 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-1040 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-1041 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-1042 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-1043 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-1044 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-1045 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

请求示例

将Pod的labels值更新为pod-test2。

```
{
  "metadata": {
    "labels": {
      "name": "pod-test2"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "annotations": {
      "kubernetes.io/availablezone": "dc1",
      "network.alpha.kubernetes.io/network": "[{\"name\":\"namespace-test-dc1-default-network\", \"interface\":\"eth0\", \"network_plane\":\"default\"}]\"",
    },
    "creationTimestamp": "2018-09-04T11:04:02Z",
    "labels": {
      "name": "pod-test2"
    },
    "name": "pod-test",
    "namespace": "namespace-test",
    "resourceVersion": "5254098",
    "selfLink": "/api/v1/namespaces/namespace-test/pods/pod-test",
    "uid": "3b99abe8-b032-11e8-9d5d-c88d83be759f"
  },
  "spec": {
    "containers": [ {
      "image": "redis:latest",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      }
    },
    "terminationMessagePath": "/dev/termination-log",
    "terminationMessagePolicy": "File"
  } ],
  "dnsPolicy": "ClusterFirst",
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "nodeName": "c0dd6256-195a-e811-90a2-10c17294fcbc",
  "restartPolicy": "Always",
  "schedulerName": "default-scheduler",
  "securityContext": { },
  "tolerations": [ {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/not-ready",
    "operator": "Exists",
    "tolerationSeconds": 300
  }, {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/unreachable",
    "operator": "Exists",
    "tolerationSeconds": 300
  } ]
}
```

```
    }
  },
  "status": {
    "conditions": [ {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-04T11:04:03Z",
      "status": "True",
      "type": "Initialized"
    }, {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-04T11:08:55Z",
      "message": "containers with unready status: [test]",
      "reason": "ContainersNotReady",
      "status": "False",
      "type": "Ready"
    }, {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-04T11:04:02Z",
      "status": "True",
      "type": "PodScheduled"
    }
  ],
  "containerStatuses": [ {
    "image": "redis:latest",
    "imageID": "",
    "lastState": {
      "terminated": {
        "containerID": "docker://f867ab7d5c68a86fc695e4d3e5f1912fdb8f98f5029ca96032b4d5d407d9a75c",
        "exitCode": 0,
        "finishedAt": "2018-09-04T11:08:33Z",
        "reason": "Completed",
        "startedAt": "2018-09-04T11:04:23Z"
      }
    },
    "name": "test",
    "ready": false,
    "restartCount": 0,
    "state": {
      "waiting": {
        "message": "rpc error: code = Unknown desc = Error response from daemon: Get https://registry-1.docker.io/v2/: net/http: request canceled while waiting for connection (Client.Timeout exceeded while awaiting headers)",
        "reason": "ErrImagePull"
      }
    }
  }
  ],
  "hostIP": "xxx.xxx.xxx.xxx",
  "phase": "Running",
  "qosClass": "Guaranteed",
  "startTime": "2018-09-04T11:04:03Z"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.7 替换 Pod

功能介绍

替换指定Pod。

其中以下字段支持更新：

- metadata.labels
- metadata.annotations
- spec.initContainers[*].image
- spec.containers[*].image
- spec.activeDeadlineSeconds
- spec.tolerations.tolerationSeconds 其余部分不支持更新。

调用方法

请参见[如何调用API](#)。

URI

PUT /api/v1/namespaces/{namespace}/pods/{name}

表 5-1046 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Pod
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1047 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1048 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-1049 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1050 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.

参数	是否必选	参数类型	描述
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.

参数	是否必选	参数类型	描述
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	是否必选	参数类型	描述
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.

参数	是否必选	参数类型	描述
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	是否必选	参数类型	描述
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceName	否	String	DeprecatedServiceAccount is a depreciated alias for ServiceAccountName. Deprecated: Use serviceName instead.

参数	是否必选	参数类型	描述
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	否	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.

参数	是否必选	参数类型	描述
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1051 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	是否必选	参数类型	描述
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1052 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1053 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1054 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1055 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1056 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1057 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1058 io.k8s.api.core.v1.PodAntiAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1059 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1060 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1061 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1062 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.
value	否	String	value is the value of the option

表 5-1063 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	是否必选	参数类型	描述
volumeDevices	否	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1064 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-1065 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1066 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1067 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1068 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1069 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1070 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-1071 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1072 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-1073 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-1074 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1075 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1076 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1077 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1078 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-1079 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1080 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1081 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1082 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-1083 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1084 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1085 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain ':'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1086 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1087 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1088 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1089 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1090 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-1091 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1092 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-1093 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	否	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	否	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1094 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1095 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1096 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1097 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1098 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1099 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-1100 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1101 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1102 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1103 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1104 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1105 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1106 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to be updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1107 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1108 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1109 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1110 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1111 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1112 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-1113 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1114 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1115 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1116 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1117 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1118 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1119 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1120 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-1121 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1122 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1123 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.

参数	是否必选	参数类型	描述
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1124 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</pre> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or</p>

参数	是否必选	参数类型	描述
			<p>suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-1125 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1126 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1127 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-1128 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-1129 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1130 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1131 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1132 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1133 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1134 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-1135 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-1136 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-1137 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1138 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1139 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1140 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1141 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1142 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1143 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1144 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmk

表 5-1145 io.k8s.api.core.v1.PodStatus

参数	是否必选	参数类型	描述
conditions	否	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	否	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	否	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	否	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	否	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	否	String	A human readable message indicating details about why the pod is in this condition.

参数	是否必选	参数类型	描述
nominatedNodeName	否	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.

参数	是否必选	参数类型	描述
phase	否	String	<p>The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values:</p> <p>Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod.</p> <p>More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase</p>
podIP	否	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.

参数	是否必选	参数类型	描述
podIPs	否	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.
qosClass	否	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	否	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	否	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-1146 io.k8s.api.core.v1.PodCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, one-word, CamelCase reason for the condition's last transition.

参数	是否必选	参数类型	描述
status	是	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	是	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-1147 io.k8s.api.core.v1.ContainerStatus

参数	是否必选	参数类型	描述
containerID	否	String	Container's ID in the format 'docker://<container_id>'.
image	是	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images
imageID	是	String	ImageID of the container's image.
lastState	否	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	是	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	是	Boolean	Specifies whether the container has passed its readiness probe.

参数	是否必选	参数类型	描述
restartCount	是	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	否	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	否	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-1148 io.k8s.api.core.v1.ContainerState

参数	是否必选	参数类型	描述
running	否	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	否	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	否	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-1149 io.k8s.api.core.v1.ContainerStateRunning

参数	是否必选	参数类型	描述
startedAt	否	String	Time at which the container was last (re-)started

表 5-1150 io.k8s.api.core.v1.ContainerStateTerminated

参数	是否必选	参数类型	描述
containerID	否	String	Container's ID in the format 'docker://<container_id>'
exitCode	是	Integer	Exit status from the last termination of the container
finishedAt	否	String	Time at which the container last terminated
message	否	String	Message regarding the last termination of the container
reason	否	String	(brief) reason from the last termination of the container
signal	否	Integer	Signal from the last termination of the container
startedAt	否	String	Time at which previous execution of the container started

表 5-1151 io.k8s.api.core.v1.ContainerStateWaiting

参数	是否必选	参数类型	描述
message	否	String	Message regarding why the container is not yet running.
reason	否	String	(brief) reason the container is not yet running.

表 5-1152 io.k8s.api.core.v1.PodIP

参数	是否必选	参数类型	描述
ip	否	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

响应参数

状态码： 200

表 5-1153 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1154 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.

参数	参数类型	描述
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerNa me	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityConte xt	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccoun t	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccoun tName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	参数类型	描述
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1155 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1156 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1157 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1158 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1159 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1160 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1161 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1162 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1163 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1164 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1165 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1166 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-1167 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1168 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-1169 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1170 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1171 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1172 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1173 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1174 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1175 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1176 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-1177 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-1178 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1179 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1180 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1181 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1182 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-1183 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1184 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1185 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1186 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-1187 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1188 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1189 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain ':'. Must not contain '*'. Must not contain '?'. Must not contain '['. Must not contain ']'. Must not contain '&'. Must not contain '='. Must not contain '". Must not contain ''.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1190 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1191 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1192 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1193 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: 1. The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1194 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-1195 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1196 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-1197 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1198 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1199 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1200 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1201 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1202 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1203 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-1204 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1205 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1206 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1207 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1208 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1209 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1210 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1211 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1212 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1213 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1214 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1215 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1216 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1217 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1218 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1219 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1220 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1221 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1222 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1223 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1224 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1225 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1226 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1227 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1228 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none">No precision is lostNo fractional digits will be emittedThe exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1229 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1230 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1231 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1232 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1233 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1234 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1235 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1236 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1237 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1238 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1239 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1240 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1241 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1242 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1243 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1244 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1245 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1246 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1247 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1248 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1249 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-1250 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-1251 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-1252 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-1253 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-1254 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-1255 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-1256 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

状态码: 201

表 5-1257 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1258 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	参数类型	描述
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1259 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1260 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1261 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1262 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1263 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1264 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1265 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1266 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1267 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1268 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1269 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1270 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-1271 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1272 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-1273 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1274 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1275 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1276 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1277 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1278 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1279 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1280 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-1281 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-1282 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1283 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1284 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1285 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1286 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-1287 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1288 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1289 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1290 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-1291 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1292 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1293 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1294 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1295 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1296 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1297 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1298 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-1299 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1300 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-1301 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1302 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1303 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1304 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1305 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1306 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1307 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-1308 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1309 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1310 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1311 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1312 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1313 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1314 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1315 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1316 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1317 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1318 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1319 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1320 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1321 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1322 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1323 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1324 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1325 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1326 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1327 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1328 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1329 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1330 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1331 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1332 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1333 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1334 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1335 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1336 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1337 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1338 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1339 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1340 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1341 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1342 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1343 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1344 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1345 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1346 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1347 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1348 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1349 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1350 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1351 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1352 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1353 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-1354 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-1355 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-1356 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-1357 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-1358 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-1359 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-1360 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

请求示例

将已创建Pod的镜像替换为"redis:latest"。

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "labels": {
      "name": "pod-test"
    },
    "name": "pod-test"
  },
  "spec": {
    "containers": [ {
      "image": "redis:latest",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "0.5",
          "memory": "1024Mi"
        },
        "requests": {
          "cpu": "0.5",
          "memory": "1024Mi"
        }
      }
    }
  ],
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "restartPolicy": "Always"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "creationTimestamp": "2018-09-04T11:04:02Z",
    "labels": {
      "name": "pod-test"
    },
    "name": "pod-test",
    "namespace": "namespace-test",
    "resourceVersion": "5253248",
    "selfLink": "/api/v1/namespaces/namespace-test/pods/pod-test",
    "uid": "3b99abe8-b032-11e8-9d5d-c88d83be759f"
  },
  "spec": {
    "containers": [ {
      "image": "redis:latest",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      }
    }
  ],
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
}
```

```
    }],
    "dnsPolicy": "ClusterFirst",
    "imagePullSecrets": [ {
      "name": "imagepull-secret"
    } ],
    "restartPolicy": "Always",
    "schedulerName": "default-scheduler",
    "securityContext": { },
    "tolerations": [ {
      "effect": "NoExecute",
      "key": "node.kubernetes.io/not-ready",
      "operator": "Exists",
      "tolerationSeconds": 300
    }, {
      "effect": "NoExecute",
      "key": "node.kubernetes.io/unreachable",
      "operator": "Exists",
      "tolerationSeconds": 300
    } ]
  },
  "status": {
    "conditions": [ {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-04T11:04:03Z",
      "status": "True",
      "type": "Initialized"
    }, {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-04T11:04:36Z",
      "status": "True",
      "type": "Ready"
    }, {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-04T11:04:02Z",
      "status": "True",
      "type": "PodScheduled"
    } ],
    "containerStatuses": [ {
      "containerID": "docker://f867ab7d5c68a86fc695e4d3e5f1912fdb8f98f5029ca96032b4d5d407d9a75c",
      "image": "redis",
      "imageID": "docker-pullable://
redis@sha256:3ab7046bd035a47aa06963d8240651d00b57e82dab07ba374ad01f84dfa1230c",
      "lastState": { },
      "name": "test",
      "ready": true,
      "restartCount": 0,
      "state": {
        "running": {
          "startedAt": "2018-09-04T11:04:23Z"
        }
      }
    } ],
    "hostIP": "xxx.xxx.xxx.xxx",
    "phase": "Running",
    "podIP": "192.168.244.170",
    "qosClass": "Guaranteed",
    "startTime": "2018-09-04T11:04:03Z"
  }
}
```

状态码

状态码	描述
200	OK
201	Created

状态码	描述
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.8 进入容器执行命令

功能介绍

进入容器执行命令。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/pods/{name}/exec

表 5-1361 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the PodExecOptions
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1362 Query 参数

参数	是否必选	参数类型	描述
command	否	String	Command is the remote command to execute. argv array. Not executed within a shell.
container	否	String	Container in which to execute the command. Defaults to only container if there is only one container in the pod.
stderr	否	Boolean	Redirect the standard error stream of the pod for this call. Defaults to true.
stdin	否	Boolean	Redirect the standard input stream of the pod for this call. Defaults to false.
stdout	否	Boolean	Redirect the standard output stream of the pod for this call. Defaults to true.
tty	否	Boolean	TTY if true indicates that a tty will be allocated for the exec call. Defaults to false.

请求参数

表 5-1363 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

无

请求示例

无

响应示例

无

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.9 进入容器执行命令

功能介绍

进入容器执行命令。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/pods/{name}/exec

表 5-1364 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the PodExecOptions
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1365 Query 参数

参数	是否必选	参数类型	描述
command	否	String	Command is the remote command to execute. argv array. Not executed within a shell.
container	否	String	Container in which to execute the command. Defaults to only container if there is only one container in the pod.
stderr	否	Boolean	Redirect the standard error stream of the pod for this call. Defaults to true.
stdin	否	Boolean	Redirect the standard input stream of the pod for this call. Defaults to false.
stdout	否	Boolean	Redirect the standard output stream of the pod for this call. Defaults to true.
tty	否	Boolean	TTY if true indicates that a tty will be allocated for the exec call. Defaults to false.

请求参数

表 5-1366 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

响应参数

无

请求示例

无

响应示例

无

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.10 查询 Pod 日志

功能介绍

查询Pod的日志。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/pods/{name}/log

表 5-1367 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Pod
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1368 Query 参数

参数	是否必选	参数类型	描述
container	否	String	The container for which to stream logs. Defaults to only container if there is one container in the pod.
follow	否	Boolean	Follow the log stream of the pod. Defaults to false.

参数	是否必选	参数类型	描述
<code>insecureSkipTLSVerifyBackend</code>	否	Boolean	<code>insecureSkipTLSVerifyBackend</code> indicates that the apiserver should not confirm the validity of the serving certificate of the backend it is connecting to. This will make the HTTPS connection between the apiserver and the backend insecure. This means the apiserver cannot verify the log data it is receiving came from the real kubelet. If the kubelet is configured to verify the apiserver's TLS credentials, it does not mean the connection to the real kubelet is vulnerable to a man in the middle attack (e.g. an attacker could not intercept the actual log data coming from the real kubelet).
<code>limitBytes</code>	否	Integer	If set, the number of bytes to read from the server before terminating the log output. This may not display a complete final line of logging, and may return slightly more or slightly less than the specified limit.
<code>pretty</code>	否	String	If 'true', then the output is pretty printed.
<code>previous</code>	否	Boolean	Return previous terminated container logs. Defaults to false.
<code>sinceSeconds</code>	否	Integer	A relative time in seconds before the current time from which to show logs. If this value precedes the time a pod was started, only logs since the pod start will be returned. If this value is in the future, no logs will be returned. Only one of <code>sinceSeconds</code> or <code>sinceTime</code> may be specified.

参数	是否必选	参数类型	描述
tailLines	否	Integer	If set, the number of lines from the end of the logs to show. If not specified, logs are shown from the creation of the container or sinceSeconds or sinceTime
timestamps	否	Boolean	If true, add an RFC3339 or RFC3339Nano timestamp at the beginning of every line of log output. Defaults to false.

请求参数

表 5-1369 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

无

请求示例

无

响应示例

无

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.11 查询 Pod 状态

功能介绍

查询Pod对象的状态。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/pods/{name}/status

表 5-1370 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Pod
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1371 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1372 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1373 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1374 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.

参数	参数类型	描述
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	参数类型	描述
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

参数	参数类型	描述
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	参数类型	描述
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1375 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1376 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1377 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1378 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1379 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1380 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1381 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1382 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1383 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1384 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1385 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1386 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-1387 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1388 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-1389 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1390 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1391 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1392 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1393 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1394 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1395 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1396 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-1397 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-1398 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1399 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1400 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1401 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1402 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-1403 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1404 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1405 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1406 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-1407 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1408 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1409 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1410 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1411 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1412 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1413 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1414 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-1415 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1416 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-1417 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1418 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1419 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1420 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	<p>FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.</p>
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	<p>Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running</p>
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	<p>GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</p>

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1421 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1422 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1423 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-1424 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1425 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1426 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1427 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1428 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1429 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1430 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1431 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1432 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1433 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1434 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1435 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1436 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1437 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1438 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1439 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1440 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1441 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1442 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1443 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1444 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1445 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1446 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1447 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1448 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1449 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1450 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1451 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1452 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1453 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1454 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1455 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1456 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1457 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1458 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1459 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1460 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1461 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1462 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1463 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1464 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1465 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1466 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1467 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1468 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1469 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-1470 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-1471 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-1472 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-1473 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-1474 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-1475 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-1476 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Pod",
  "metadata": {
    "annotations": {
      "cri.cci.io/container-type": "secure-container",
      "kubernetes.io/availablezone": "dc1",
      "network.alpha.kubernetes.io/network": "[{"name": "namespace-test-dc1-default-network", "interface": "eth0", "network_plane": "default"}]"
    },
    "creationTimestamp": "2018-09-03T12:26:12Z",
    "labels": {
      "name": "pod-test"
    },
    "name": "pod-test",
    "namespace": "namespace-test",
    "resourceVersion": "5030610",
    "selfLink": "/api/v1/namespaces/namespace-test/pods/pod-test",
    "uid": "8b985a27-af74-11e8-9d5d-c88d83be759f"
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "Always",
      "name": "test",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      }
    }
  ],
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
},
  "dnsPolicy": "ClusterFirst",
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "nodeName": "c0dd6256-195a-e811-90a2-10c17294fcbc",
  "restartPolicy": "Always",
  "schedulerName": "default-scheduler",
  "securityContext": { },
  "tolerations": [ {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/not-ready",
    "operator": "Exists",
    "tolerationSeconds": 300
  }, {
    "effect": "NoExecute",
    "key": "node.kubernetes.io/unreachable",
    "operator": "Exists",
    "tolerationSeconds": 300
  } ]
},
  "status": {
    "conditions": [ {
      "lastProbeTime": null,
      "lastTransitionTime": "2018-09-03T12:26:12Z",
      "status": "True",
```



```

    "type" : "Initialized"
  }, {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2018-09-03T12:26:16Z",
    "status" : "True",
    "type" : "Ready"
  }, {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2018-09-03T12:26:12Z",
    "status" : "True",
    "type" : "PodScheduled"
  } ],
  "containerStatuses" : [ {
    "containerID" : "docker://aee55d8dedb8371f96aa5d5116f69a53bf1cb23afe1802567c24081514d3b048",
    "image" : "redis",
    "imageID" : "docker-pullable://
redis@sha256:3ab7046bd035a47aa06963d8240651d00b57e82dab07ba374ad01f84dfa1230c",
    "lastState" : { },
    "name" : "test",
    "ready" : true,
    "restartCount" : 0,
    "state" : {
      "running" : {
        "startedAt" : "2018-09-03T12:26:16Z"
      }
    }
  } ],
  "phase" : "Running",
  "podIP" : "192.168.245.185",
  "qosClass" : "Guaranteed",
  "startTime" : "2018-09-03T12:26:12Z"
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.2.12 查询用户所有的 Pods

功能介绍

该API用于获取一个Pod列表。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/pods

表 5-1477 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 5-1478 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1479 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Pod objects	List of pods. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-1480 io.k8s.api.core.v1.Pod

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.PodStatus object	Most recently observed status of the pod. This data may not be up to date. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1481 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	参数类型	描述
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.

参数	参数类型	描述
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1482 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1483 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1484 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1485 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1486 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1487 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1488 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1489 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1490 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1491 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1492 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1493 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-1494 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1495 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-1496 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1497 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1498 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1499 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1500 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1501 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1502 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1503 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-1504 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-1505 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1506 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1507 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1508 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1509 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-1510 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1511 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1512 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1513 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-1514 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1515 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1516 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1517 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1518 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1519 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1520 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1521 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-1522 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1523 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-1524 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1525 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1526 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1527 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1528 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1529 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1530 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-1531 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1532 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1533 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1534 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1535 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1536 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1537 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1538 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1539 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1540 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1541 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1542 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1543 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1544 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1545 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1546 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1547 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1548 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1549 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1550 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1551 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1552 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1553 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1554 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1555 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1556 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1557 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1558 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1559 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1560 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1561 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1562 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1563 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1564 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1565 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1566 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1567 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1568 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1569 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1570 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1571 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1572 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1573 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1574 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1575 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1576 io.k8s.api.core.v1.PodStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.PodCondition objects	Current service state of pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
containerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per container in the manifest. Each entry is currently the output of <i>docker inspect</i> . More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
ephemeralContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	Status for any ephemeral containers that have run in this pod. This field is alpha-level and is only populated by servers that enable the EphemeralContainers feature.
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
initContainerStatuses	Array of io.k8s.api.core.v1.ContainerStatus objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-and-container-status
message	String	A human readable message indicating details about why the pod is in this condition.

参数	参数类型	描述
nominatedNodeName	String	nominatedNodeName is set only when this pod preempts other pods on the node, but it cannot be scheduled right away as preemption victims receive their graceful termination periods. This field does not guarantee that the pod will be scheduled on this node. Scheduler may decide to place the pod elsewhere if other nodes become available sooner. Scheduler may also decide to give the resources on this node to a higher priority pod that is created after preemption. As a result, this field may be different than PodSpec.nodeName when the pod is scheduled.
phase	String	The phase of a Pod is a simple, high-level summary of where the Pod is in its lifecycle. The conditions array, the reason and message fields, and the individual container status arrays contain more detail about the pod's status. There are five possible phase values: Pending: The pod has been accepted by the Kubernetes system, but one or more of the container images has not been created. This includes time before being scheduled as well as time spent downloading images over the network, which could take a while. Running: The pod has been bound to a node, and all of the containers have been created. At least one container is still running, or is in the process of starting or restarting. Succeeded: All containers in the pod have terminated in success, and will not be restarted. Failed: All containers in the pod have terminated, and at least one container has terminated in failure. The container either exited with non-zero status or was terminated by the system. Unknown: For some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-phase
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
podIPs	Array of io.k8s.api.core.v1.PodIP objects	podIPs holds the IP addresses allocated to the pod. If this field is specified, the 0th entry must match the podIP field. Pods may be allocated at most 1 value for each of IPv4 and IPv6. This list is empty if no IPs have been allocated yet.

参数	参数类型	描述
qosClass	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements See PodQOSClass type for available QOS classes More info: https://git.k8s.io/community/contributors/design-proposals/node/resource-qos.md
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'Evicted'
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

表 5-1577 io.k8s.api.core.v1.PodCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
status	String	Status is the status of the condition. Can be True, False, Unknown. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions
type	String	Type is the type of the condition. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#pod-conditions

表 5-1578 io.k8s.api.core.v1.ContainerStatus

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'.
image	String	The image the container is running. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imageID	String	ImageID of the container's image.
lastState	io.k8s.api.core.v1.ContainerState object	Details about the container's last termination condition.
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
ready	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. But those containers are subject to garbage collection. This value will get capped at 5 by GC.
started	Boolean	Specifies whether the container has passed its startup probe. Initialized as false, becomes true after startupProbe is considered successful. Resets to false when the container is restarted, or if kubelet loses state temporarily. Is always true when no startupProbe is defined.
state	io.k8s.api.core.v1.ContainerState object	Details about the container's current condition.

表 5-1579 io.k8s.api.core.v1.ContainerState

参数	参数类型	描述
running	io.k8s.api.core.v1.ContainerStateRunning object	Details about a running container
terminated	io.k8s.api.core.v1.ContainerStateTerminated object	Details about a terminated container
waiting	io.k8s.api.core.v1.ContainerStateWaiting object	Details about a waiting container

表 5-1580 io.k8s.api.core.v1.ContainerStateRunning

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started

表 5-1581 io.k8s.api.core.v1.ContainerStateTerminated

参数	参数类型	描述
containerID	String	Container's ID in the format 'docker://<container_id>'
exitCode	Integer	Exit status from the last termination of the container
finishedAt	String	Time at which the container last terminated
message	String	Message regarding the last termination of the container
reason	String	(brief) reason from the last termination of the container
signal	Integer	Signal from the last termination of the container
startedAt	String	Time at which previous execution of the container started

表 5-1582 io.k8s.api.core.v1.ContainerStateWaiting

参数	参数类型	描述
message	String	Message regarding why the container is not yet running.
reason	String	(brief) reason the container is not yet running.

表 5-1583 io.k8s.api.core.v1.PodIP

参数	参数类型	描述
ip	String	ip is an IP address (IPv4 or IPv6) assigned to the pod

表 5-1584 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "annotations": {
        "kubernetes.io/created-by": "{\"kind\":\"SerializedReference\",\"apiVersion\":\"v1\",\"reference\":{\"kind\":\"ReplicationController\",\"namespace\":\"default\",\"name\":\"fdsfsd\",\"uid\":\"7539c329-57c4-11e7-afb7-fa163e218692\"},\"apiVersion\":\"v1\",\"resourceVersion\":\"956153\"}}",
        "kubernetes.io/limit-ranger": "LimitRanger plugin set: cpu request for container container01"
      },
      "creationTimestamp": "2017-06-23T03:31:35Z",
      "generateName": "fdsfsd-",
      "labels": {
        "cce/appgroup": "gfsad",
        "name": "fdsfsd"
      },
      "name": "fdsfsd-ddnft",
      "namespace": "default",
      "ownerReferences": [ {
        "apiVersion": "v1",
        "controller": true,
        "kind": "ReplicationController",
        "name": "fdsfsd",
        "uid": "7539c329-57c4-11e7-afb7-fa163e218692"
      } ],
      "resourceVersion": "1449034",
      "selfLink": "/api/v1/namespaces/default/pods/fdsfsd-ddnft",
      "uid": "753a45bc-57c4-11e7-afb7-fa163e218692"
    },
    "spec": {
      "containers": [ {
        "image": "10.154.52.159:443/test/apache-php:latest",
        "imagePullPolicy": "Always",
        "name": "container01",
        "ports": [ {
          "containerPort": 80,
          "protocol": "TCP"
        } ],
        "resources": {
          "requests": {
            "cpu": "100m"
          }
        }
      } ],
      "terminationMessagePath": "/dev/termination-log",
      "volumeMounts": [ {
        "mountPath": "/var/run/secrets/kubernetes.io/serviceaccount",
        "name": "default-token-863rh",
        "readOnly": true
      } ]
    },
    "dnsPolicy": "ClusterFirst",
    "imagePullSecrets": [ {
      "name": "myregistry"
    } ],
    "nodeName": "192.168.12.187",
    "restartPolicy": "Always",
    "securityContext": { },
    "serviceAccount": "default",
  } ]
}
```

```
"serviceAccountName" : "default",
"volumes" : [ {
  "name" : "default-token-863rh",
  "secret" : {
    "defaultMode" : 420,
    "secretName" : "default-token-863rh"
  }
} ]
},
"status" : {
  "conditions" : [ {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2017-06-23T03:31:36Z",
    "status" : "True",
    "type" : "Initialized"
  }, {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2017-06-28T06:34:51Z",
    "status" : "True",
    "type" : "Ready"
  }, {
    "lastProbeTime" : null,
    "lastTransitionTime" : "2017-06-23T03:31:35Z",
    "status" : "True",
    "type" : "PodScheduled"
  } ],
  "containerStatuses" : [ {
    "containerID" : "docker://f3daa802f753d123fe66b2cba2e917725702f8d446c17541822a68f9d2414c6d",
    "image" : "10.154.52.159:443/test/apache-php:latest",
    "imageID" : "docker://
sha256:2e233ad9329bd7f65572dd6acb1a03e8839c36abfdb1d1f9012d84d13cecf9fc",
    "lastState" : { },
    "name" : "container01",
    "ready" : true,
    "restartCount" : 0,
    "state" : {
      "running" : {
        "startedAt" : "2017-06-28T06:34:46Z"
      }
    }
  } ],
  "hostIP" : "192.168.12.187",
  "phase" : "Running",
  "podIP" : "172.16.56.4",
  "startTime" : "2017-06-23T03:31:36Z"
}
} ],
"kind" : "PodList",
"metadata" : {
  "resourceVersion" : "1550321",
  "selfLink" : "/api/v1/pods"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.3 StorageClass

5.3.1 查询/apis/storage.k8s.io/v1 版本的所有 storage class

功能介绍

list or watch objects of kind StorageClass

调用方法

请参见[如何调用API](#)。

URI

GET /apis/storage.k8s.io/v1/storageclasses

表 5-1585 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1586 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1587 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.storage.v1.StorageClass objects	Items is the list of StorageClasses
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-1588 io.k8s.api.storage.v1.StorageClass

参数	参数类型	描述
allowVolumeExpansion	Boolean	AllowVolumeExpansion shows whether the storage class allow volume expand
allowedTopologies	Array of io.k8s.api.core.v1.TopologySelectorTerm objects	Restrict the node topologies where volumes can be dynamically provisioned. Each volume plugin defines its own supported topology specifications. An empty TopologySelectorTerm list means there is no topology restriction. This field is only honored by servers that enable the VolumeScheduling feature.

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
mountOptions	Array of strings	Dynamically provisioned PersistentVolumes of this storage class are created with these mountOptions, e.g. ["ro", "soft"]. Not validated <ul style="list-style-type: none">mount of the PVs will simply fail if one is invalid.
parameters	Map<String,String>	Parameters holds the parameters for the provisioner that should create volumes of this storage class.
provisioner	String	Provisioner indicates the type of the provisioner.
reclaimPolicy	String	Dynamically provisioned PersistentVolumes of this storage class are created with this reclaimPolicy. Defaults to Delete.
volumeBindingMode	String	VolumeBindingMode indicates how PersistentVolumeClaims should be provisioned and bound. When unset, VolumeBindingImmediate is used. This field is only honored by servers that enable the VolumeScheduling feature.

表 5-1589 io.k8s.api.core.v1.TopologySelectorTerm

参数	参数类型	描述
matchLabelExpressions	Array of io.k8s.api.core.v1.TopologySelectorLabelRequirement objects	A list of topology selector requirements by labels.

表 5-1590 io.k8s.api.core.v1.TopologySelectorLabelRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
values	Array of strings	An array of string values. One value must match the label to be selected. Each entry in Values is ORed.

表 5-1591 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1592 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1593 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1594 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
[ {
  "allowVolumeExpansion" : false,
  "metadata" : {
    "creationTimestamp" : "2022-07-26T04:58:34Z",
    "name" : "csi-disk-sas",
    "resourceVersion" : "289546",
    "selfLink" : "/apis/storage.k8s.io/v1/storageclasses/csi-disk-sas",
    "uid" : "f47fe78e-1677-4ac4-b02c-109cee63f34a"
  },
  "parameters" : {
    "csi.storage.k8s.io/csi-driver-name" : "disk.csi.everest.io",
    "csi.storage.k8s.io/fstype" : "ext4",
    "everest.io/disk-volume-type" : "SAS",
    "everest.io/passthrough" : "true"
  },
  "provisioner" : "everest-csi-provisioner",
  "reclaimPolicy" : "Delete",
  "volumeBindingMode" : "Immediate"
}, {
  "allowVolumeExpansion" : false,
  "metadata" : {
    "creationTimestamp" : "2022-07-26T04:58:34Z",
    "name" : "csi-disk-sata",
    "resourceVersion" : "289547",
    "selfLink" : "/apis/storage.k8s.io/v1/storageclasses/csi-disk-sata",
    "uid" : "3df529ff-5300-4bd8-a292-34ed5ddd80cd"
  },
  "parameters" : {
    "csi.storage.k8s.io/csi-driver-name" : "disk.csi.everest.io",
    "csi.storage.k8s.io/fstype" : "ext4",
    "everest.io/disk-volume-type" : "SATA",
    "everest.io/passthrough" : "true"
  },
  "provisioner" : "everest-csi-provisioner",
  "reclaimPolicy" : "Delete",
  "volumeBindingMode" : "Immediate"
} ]
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.3.2 查询指定的 storage class

功能介绍

read the specified StorageClass

调用方法

请参见[如何调用API](#)。

URI

GET /apis/storage.k8s.io/v1/storageclasses/{name}

表 5-1595 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the StorageClass

表 5-1596 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.

参数	是否必选	参数类型	描述
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1597 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1598 响应 Body 参数

参数	参数类型	描述
allowVolumeExpansion	Boolean	AllowVolumeExpansion shows whether the storage class allow volume expand
allowedTopologies	Array of io.k8s.api.core.v1.TopologySelectorTerm objects	Restrict the node topologies where volumes can be dynamically provisioned. Each volume plugin defines its own supported topology specifications. An empty TopologySelectorTerm list means there is no topology restriction. This field is only honored by servers that enable the VolumeScheduling feature.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
mountOptions	Array of strings	Dynamically provisioned PersistentVolumes of this storage class are created with these mountOptions, e.g. ["ro", "soft"]. Not validated <ul style="list-style-type: none">mount of the PVs will simply fail if one is invalid.
parameters	Map<String,String>	Parameters holds the parameters for the provisioner that should create volumes of this storage class.
provisioner	String	Provisioner indicates the type of the provisioner.
reclaimPolicy	String	Dynamically provisioned PersistentVolumes of this storage class are created with this reclaimPolicy. Defaults to Delete.
volumeBindingMode	String	VolumeBindingMode indicates how PersistentVolumeClaims should be provisioned and bound. When unset, VolumeBindingImmediate is used. This field is only honored by servers that enable the VolumeScheduling feature.

表 5-1599 io.k8s.api.core.v1.TopologySelectorTerm

参数	参数类型	描述
matchLabelExpressions	Array of io.k8s.api.core.v1.TopologySelectorLabelRequirement objects	A list of topology selector requirements by labels.

表 5-1600 io.k8s.api.core.v1.TopologySelectorLabelRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
values	Array of strings	An array of string values. One value must match the label to be selected. Each entry in Values is ORed.

表 5-1601 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1602 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1603 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

无

响应示例

状态码: 200

OK

```
{  
  "allowVolumeExpansion" : false,  
}
```



```
"metadata": {
  "creationTimestamp": "2022-07-26T04:58:34Z",
  "name": "storage-test",
  "resourceVersion": "289547",
  "selfLink": "/apis/storage.k8s.io/v1/storageclasses/storage-test",
  "uid": "3df529ff-5300-4bd8-a292-34ed5ddd80cd"
},
"parameters": {
  "csi.storage.k8s.io/csi-driver-name": "disk.csi.everest.io",
  "csi.storage.k8s.io/fstype": "ext4",
  "everest.io/disk-volume-type": "SATA",
  "everest.io/passthrough": "true"
},
"provisioner": "everest-csi-provisioner",
"reclaimPolicy": "Delete",
"volumeBindingMode": "Immediate"
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4 Service

5.4.1 查询指定 namespace 下的 Services

功能介绍

查询Namespace下所有Service的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/services

表 5-1604 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1605 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1606 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1607 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Service objects	List of services
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-1608 io.k8s.api.core.v1.Service

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1609 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1610 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1611 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1612 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1613 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1614 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1615 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1616 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1617 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1618 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

表 5-1619 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2018-09-03T11:20:56Z",
      "labels": {
        "addonmanager.kubernetes.io/mode": "Reconcile",
        "app": "kube-dns",
        "kubernetes.io/cluster-service": "true",
        "kubernetes.io/name": "KubeDNS"
      },
      "name": "kube-dns",
      "namespace": "namespace-test",
      "resourceVersion": "5016792",
      "selfLink": "/api/v1/namespaces/namespace-test/services/kube-dns",
      "uid": "6d356f18-af6b-11e8-b6ef-f898ef6c78b4"
    },
    "spec": {
      "clusterIP": "10.247.189.43",
      "ports": [ {
        "name": "dns",
        "port": 53,
        "protocol": "UDP",
        "targetPort": 5353
      }, {
        "name": "dns-tcp",
        "port": 53,
        "protocol": "TCP",
        "targetPort": 5353
      } ],
    }
  } ],
}
```

```
"selector" : {
  "app" : "kube-dns"
},
"sessionAffinity" : "None",
"type" : "ClusterIP"
},
"status" : {
  "loadBalancer" : {}
}
}, {
"metadata" : {
  "creationTimestamp" : "2018-09-04T00:45:36Z",
  "labels" : {
    "app" : "redis"
  },
  "name" : "redis",
  "namespace" : "namespace-test",
  "resourceVersion" : "5146412",
  "selfLink" : "/api/v1/namespaces/namespace-test/services/redis",
  "uid" : "d6a1ce79-afdb-11e8-b6ef-f898ef6c78b4"
},
"spec" : {
  "clusterIP" : "10.247.212.210",
  "ports" : [ {
    "name" : "service0",
    "port" : 8080,
    "protocol" : "TCP",
    "targetPort" : 80
  } ],
  "selector" : {
    "app" : "redis"
  },
  "sessionAffinity" : "None",
  "type" : "ClusterIP"
},
"status" : {
  "loadBalancer" : {}
}
}],
"kind" : "ServiceList",
"metadata" : {
  "resourceVersion" : "5147871",
  "selfLink" : "/api/v1/namespaces/namespace-test/services"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict

状态码	描述
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4.2 创建 Service

功能介绍

创建一个Service。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/services

表 5-1620 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1621 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1622 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-1623 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	是否必选	参数类型	描述
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1624 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	是否必选	参数类型	描述
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1625 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1626 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1627 io.k8s.api.core.v1.ServiceSpec

参数	是否必选	参数类型	描述
clusterIP	否	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	否	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	否	String	externalName is the external reference that kubernetes or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.

参数	是否必选	参数类型	描述
externalTrafficPolicy	否	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	否	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	是否必选	参数类型	描述
ipFamily	否	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)

参数	是否必选	参数类型	描述
loadBalancerIP	否	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	否	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	否	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	是否必选	参数类型	描述
publishNotReadyAddresses	否	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	否	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	否	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	是否必选	参数类型	描述
sessionAffinityConfig	否	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.
topologyKeys	否	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.

参数	是否必选	参数类型	描述
type	否	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1628 io.k8s.api.core.v1.ServicePort

参数	是否必选	参数类型	描述
appProtocol	否	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Un-prefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.
name	否	String	The name of this port within the service. This must be a <code>DNS_LABEL</code> . All ports within a <code>ServiceSpec</code> must have unique names. When considering the endpoints for a <code>Service</code> , this must match the 'name' field in the <code>EndpointPort</code> . Optional if only one <code>ServicePort</code> is defined on this service.
nodePort	否	Integer	The port on each node on which this service is exposed when <code>type=NodePort</code> or <code>LoadBalancer</code> . Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the <code>ServiceType</code> of this <code>Service</code> requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	是	Integer	The port that will be exposed by this service.
protocol	否	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.

参数	是否必选	参数类型	描述
targetPort	否	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1629 io.k8s.api.core.v1.SessionAffinityConfig

参数	是否必选	参数类型	描述
clientIP	否	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1630 io.k8s.api.core.v1.ClientIPConfig

参数	是否必选	参数类型	描述
timeoutSeconds	否	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1631 io.k8s.api.core.v1.ServiceStatus

参数	是否必选	参数类型	描述
loadBalancer	否	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1632 io.k8s.api.core.v1.LoadBalancerStatus

参数	是否必选	参数类型	描述
ingress	否	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1633 io.k8s.api.core.v1.LoadBalancerIngress

参数	是否必选	参数类型	描述
hostname	否	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	否	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

响应参数

状态码： 200

表 5-1634 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1635 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1636 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1637 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1638 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1639 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1640 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1641 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1642 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1643 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1644 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

状态码： 201

表 5-1645 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1646 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1647 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1648 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1649 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1650 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1651 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1652 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1653 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1654 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1655 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

状态码： 202

表 5-1656 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1657 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1658 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1659 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1660 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1661 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1662 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1663 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1664 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1665 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1666 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

- 创建ClusterIP类型的Service，公开的端口为80，端口的IP协议为"TCP"。

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "labels": {
      "app": "redis"
    },
    "name": "redis"
  },
  "spec": {
    "ports": [ {
      "name": "service0",
      "port": 8080,
      "protocol": "TCP",
      "targetPort": 80
    } ],
    "selector": {
      "app": "redis"
    },
    "type": "ClusterIP"
  }
}
```

- 创建LoadBalancer类型的Service，指定elb实例ID、项目ID和账号ID。
LoadBalancer类型Service需要在metadata.annotations自定义中添加elb实例ID（kubernetes.io/elb.id）、项目ID（tenant.kubernetes.io/project-id）和账号ID（tenant.kubernetes.io/domain-id）

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "77e6246c-a091-xxxx-xxxx-789baa571280",
      "tenant.kubernetes.io/domain-id": "65382xxxxxxxxxxxxxxxxxxxxe684b",
      "tenant.kubernetes.io/project-id": "a9cab8xxxxxxxxxxxxxxxx41c0aeb"
    },
    "name": "nginx"
  },
  "spec": {
    "ports": [ {
      "name": "service0",
      "port": 8080,
      "protocol": "TCP",
      "targetPort": 80
    } ],
    "selector": {
      "app": "nginx"
    },
    "type": "LoadBalancer"
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "creationTimestamp": "2018-09-04T00:45:36Z",
    "labels": {
      "app": "redis"
    },
    "name": "redis",
    "namespace": "namespace-test",
    "resourceVersion": "5146412",
    "selfLink": "/api/v1/namespaces/namespace-test/services/redis",
    "uid": "d6a1ce79-afdb-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "clusterIP": "10.247.212.210",
    "ports": [ {
      "name": "service0",
      "port": 8080,
      "protocol": "TCP",
      "targetPort": 80
    } ],
    "selector": {
      "app": "redis"
    },
    "sessionAffinity": "None",
    "type": "ClusterIP"
  },
  "status": {
    "loadBalancer": { }
  }
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4.3 删除 Service

功能介绍

删除Service。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/services/{name}

表 5-1667 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Service
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1668 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1669 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-1670 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-1671 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-1672 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.

参数	参数类型	描述
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1673 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.

参数	参数类型	描述
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1674 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-1675 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-1676 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1677 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1678 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-1679 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "code": 200,
  "kind": "Status",
  "metadata": { },
  "status": "Success"
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4.4 查询 Service

功能介绍

查询Service的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/services/{name}

表 5-1680 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Service
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1681 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.

参数	是否必选	参数类型	描述
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1682 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1683 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1684 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1685 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1686 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1687 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1688 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1689 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1690 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1691 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1692 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1693 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "creationTimestamp": "2018-09-04T00:45:36Z",
    "labels": {
      "app": "redis"
    }
  },
  "name": "redis",
  "namespace": "namespace-test",
  "resourceVersion": "5146412",
  "selfLink": "/api/v1/namespaces/namespace-test/services/redis",
  "uid": "d6a1ce79-afdb-11e8-b6ef-f898ef6c78b4"
},
"spec": {
  "clusterIP": "10.247.212.210",
  "ports": [ {
    "name": "service0",
    "port": 8080,
    "protocol": "TCP",
    "targetPort": 80
  } ],
  "selector": {
    "app": "redis"
  },
  "sessionAffinity": "None",
  "type": "ClusterIP"
},
"status": {
  "loadBalancer": { }
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4.5 更新 Service

功能介绍

更新Service。

调用方法

请参见[如何调用API](#)。

URI

PATCH /api/v1/namespaces/{namespace}/services/{name}

表 5-1694 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Service
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1695 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1696 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-1697 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-1698 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1699 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1700 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1701 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1702 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1703 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1704 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1705 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1706 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1707 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1708 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

更新Service中的labels值为"some-key" : "some-value"。

```
{
  "metadata": {
    "labels": {
      "some-key": "some-value"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "creationTimestamp": "2022-09-08T07:14:35Z",
    "labels": {
      "app": "service-test",
      "some-key": "some-value"
    },
    "name": "service-test",
    "namespace": "namespace-test",
    "resourceVersion": "43938742",
    "selfLink": "/api/v1/namespaces/namespace-test/services/service-test",
    "uid": "4397866e-7583-42e3-b428-2381045e1513"
  },
  "spec": {
    "clusterIP": "10.247.175.81",
    "ports": [ {
      "name": "service0",
      "port": 80,
      "protocol": "TCP",
      "targetPort": 80
    } ],
    "selector": {
      "app": "service-test"
    },
    "sessionAffinity": "None",
    "type": "ClusterIP"
  },
  "status": {
    "loadBalancer": { }
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4.6 替换 Service

功能介绍

替换Service。

调用方法

请参见[如何调用API](#)。

URI

PUT /api/v1/namespaces/{namespace}/services/{name}

表 5-1709 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Service
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1710 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1711 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-1712 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1713 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1714 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1715 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1716 io.k8s.api.core.v1.ServiceSpec

参数	是否必选	参数类型	描述
clusterIP	否	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	否	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	否	String	externalName is the external reference that kubernetes or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.

参数	是否必选	参数类型	描述
externalTrafficPolicy	否	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	否	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	是否必选	参数类型	描述
ipFamily	否	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)

参数	是否必选	参数类型	描述
loadBalancerIP	否	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	否	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	否	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	是否必选	参数类型	描述
publishNotReadyAddresses	否	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	否	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	否	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	是否必选	参数类型	描述
sessionAffinityConfig	否	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.
topologyKeys	否	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.

参数	是否必选	参数类型	描述
type	否	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1717 io.k8s.api.core.v1.ServicePort

参数	是否必选	参数类型	描述
appProtocol	否	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Un-prefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.
name	否	String	The name of this port within the service. This must be a <code>DNS_LABEL</code> . All ports within a <code>ServiceSpec</code> must have unique names. When considering the endpoints for a <code>Service</code> , this must match the 'name' field in the <code>EndpointPort</code> . Optional if only one <code>ServicePort</code> is defined on this service.
nodePort	否	Integer	The port on each node on which this service is exposed when <code>type=NodePort</code> or <code>LoadBalancer</code> . Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the <code>ServiceType</code> of this <code>Service</code> requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	是	Integer	The port that will be exposed by this service.
protocol	否	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.

参数	是否必选	参数类型	描述
targetPort	否	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1718 io.k8s.api.core.v1.SessionAffinityConfig

参数	是否必选	参数类型	描述
clientIP	否	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1719 io.k8s.api.core.v1.ClientIPConfig

参数	是否必选	参数类型	描述
timeoutSeconds	否	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1720 io.k8s.api.core.v1.ServiceStatus

参数	是否必选	参数类型	描述
loadBalancer	否	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1721 io.k8s.api.core.v1.LoadBalancerStatus

参数	是否必选	参数类型	描述
ingress	否	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1722 io.k8s.api.core.v1.LoadBalancerIngress

参数	是否必选	参数类型	描述
hostname	否	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	否	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

响应参数

状态码： 200

表 5-1723 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1724 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1725 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1726 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1727 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1728 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1729 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1730 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1731 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1732 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1733 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

状态码： 201

表 5-1734 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1735 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1736 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1737 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1738 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1739 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1740 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1741 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1742 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1743 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1744 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

将已创建Service的名称替换为“service-test”。

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.class": "dnat",
      "kubernetes.io/natgateway.id": "4ed25dd6-1887-439f-aef6-b1e3f2b57b5c",
      "tenant.kubernetes.io/domain-id": "08a2c8ef8180d4150ff5c0012463ee60",
      "tenant.kubernetes.io/project-id": "08a2c8ef8d80d4152ff8c001d0281c03"
    },
    "creationTimestamp": "2022-09-06T06:28:09Z",
    "finalizers": [ "service.kubernetes.io/load-balancer-cleanup" ],
    "labels": {
      "app": "service-test"
    },
    "name": "service-test",
    "namespace": "namespace-test",
    "resourceVersion": "41521168",
    "selfLink": "/api/v1/namespaces/namespace-test/services/service-test",
    "uid": "7dfc42ef-f938-401c-b44a-1f7bd79b3fcb"
  },
  "spec": {
    "clusterIP": "10.247.64.172",
    "externalTrafficPolicy": "Cluster",
    "loadBalancerIP": "100.93.1.98",
    "ports": [ {
      "name": "service0",
      "nodePort": 31966,
      "port": 30157,
      "protocol": "TCP",
      "targetPort": 80
    } ],
    "selector": {
      "app": "service-test"
    },
    "sessionAffinity": "None",
    "type": "LoadBalancer"
  },
  "status": {
    "loadBalancer": { }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
```

```
"annotations": {
  "kubernetes.io/elb.class": "dnat",
  "kubernetes.io/natgateway.id": "4ed25dd6-1887-439f-ae6f-b1e3f2b57b5c",
  "tenant.kubernetes.io/domain-id": "08a2c8ef8180d4150ff5c0012463ee60",
  "tenant.kubernetes.io/project-id": "08a2c8ef8d80d4152ff8c001d0281c03"
},
"creationTimestamp": "2022-09-06T06:28:09Z",
"finalizers": [ "service.kubernetes.io/load-balancer-cleanup" ],
"labels": {
  "app": "service-test"
},
"name": "service-test",
"namespace": "namespace-test",
"resourceVersion": "41521168",
"selfLink": "/api/v1/namespaces/namespace-test/services/service-test",
"uid": "7dfc42ef-f938-401c-b44a-1f7bd79b3fcb"
},
"spec": {
  "clusterIP": "10.247.64.172",
  "externalTrafficPolicy": "Cluster",
  "loadBalancerIP": "100.93.1.98",
  "ports": [ {
    "name": "service0",
    "nodePort": 31966,
    "port": 30157,
    "protocol": "TCP",
    "targetPort": 80
  } ],
  "selector": {
    "app": "service-test"
  },
  "sessionAffinity": "None",
  "type": "LoadBalancer"
},
"status": {
  "loadBalancer": { }
}
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid

状态码	描述
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.4.7 查询 service 状态

功能介绍

查询指定的Service的状态。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/services/{name}/status

表 5-1745 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Service
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1746 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1747 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1748 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ServiceSpec object	Spec defines the behavior of a service. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ServiceStatus object	Most recently observed status of the service. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1749 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1750 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1751 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1752 io.k8s.api.core.v1.ServiceSpec

参数	参数类型	描述
clusterIP	String	clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field can not be changed through updates. Valid values are "None", empty string (""), or a valid IP address. "None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
externalIPs	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalName	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid RFC-1123 hostname (https://tools.ietf.org/html/rfc1123) and requires Type to be ExternalName.
externalTrafficPolicy	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading. "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	参数类型	描述
ipFamily	String	ipFamily specifies whether this Service has a preference for a particular IP family (e.g. IPv4 vs. IPv6) when the IPv6DualStack feature gate is enabled. In a dual-stack cluster, you can specify ipFamily when creating a ClusterIP Service to determine whether the controller will allocate an IPv4 or IPv6 IP for it, and you can specify ipFamily when creating a headless Service to determine whether it will have IPv4 or IPv6 Endpoints. In either case, if you do not specify an ipFamily explicitly, it will default to the cluster's primary IP family. This field is part of an alpha feature, and you should not make any assumptions about its semantics other than those described above. In particular, you should not assume that it can (or cannot) be changed after creation time; that it can only have the values "IPv4" and "IPv6"; or that its current value on a given Service correctly reflects the current state of that Service. (For ClusterIP Services, look at clusterIP to see if the Service is IPv4 or IPv6. For headless Services, look at the endpoints, which may be dual-stack in the future. For ExternalName Services, ipFamily has no meaning, but it may be set to an irrelevant value anyway.)
loadBalancerIP	String	Only applies to Service Type: LoadBalancer. LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	Array of strings	If specified and supported by the platform, this will restrict traffic through the cloud-provider load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature." More info: https://kubernetes.io/docs/tasks/access-application-cluster/configure-cloud-provider-firewall/
ports	Array of io.k8s.api.core.v1.ServicePort objects	The list of ports that are exposed by this service. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies

参数	参数类型	描述
publishNotReadyAddresses	Boolean	publishNotReadyAddresses indicates that any agent which deals with endpoints for this Service should disregard any indications of ready/not-ready. The primary use case for setting this field is for a StatefulSet's Headless Service to propagate SRV DNS records for its Pods for the purpose of peer discovery. The Kubernetes controllers that generate Endpoints and EndpointSlice resources for Services interpret this to mean that all endpoints are considered "ready" even if the Pods themselves are not. Agents which consume only Kubernetes generated endpoints through the Endpoints or EndpointSlice resources can safely assume this behavior.
selector	Map<String,String>	Route service traffic to pods with label keys and values matching this selector. If empty or not present, the service is assumed to have an external process managing its endpoints, which Kubernetes will not modify. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName. More info: https://kubernetes.io/docs/concepts/services-networking/service/
sessionAffinity	String	Supports "ClientIP" and "None". Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None. More info: https://kubernetes.io/docs/concepts/services-networking/service/#virtual-ips-and-service-proxies
sessionAffinityConfig	io.k8s.api.core.v1.SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

参数	参数类型	描述
topologyKeys	Array of strings	topologyKeys is a preference-order list of topology keys which implementations of services should use to preferentially sort endpoints when accessing this Service, it can not be used at the same time as externalTrafficPolicy=Local. Topology keys must be valid label keys and at most 16 keys may be specified. Endpoints are chosen based on the first topology key with available backends. If this field is specified and all entries have no backends that match the topology of the client, the service has no backends for that client and connections should fail. The special value "*" may be used to mean "any topology". This catch-all value, if used, only makes sense as the last value in the list. If this is not specified or empty, no topology constraints will be applied.
type	String	type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP, NodePort, and LoadBalancer. "ExternalName" maps to the specified externalName. "ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP. "NodePort" builds on ClusterIP and allocates a port on every node which routes to the clusterIP. "LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP. More info: https://kubernetes.io/docs/concepts/services-networking/service/#publishing-services-service-types

表 5-1753 io.k8s.api.core.v1.ServicePort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	String	The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. When considering the endpoints for a Service, this must match the 'name' field in the EndpointPort. Optional if only one ServicePort is defined on this service.
nodePort	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. More info: https://kubernetes.io/docs/concepts/services-networking/service/#type-nodeport
port	Integer	The port that will be exposed by this service.
protocol	String	The IP protocol for this port. Supports "TCP", "UDP", and "SCTP". Default is TCP.
targetPort	String	Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of the 'port' field is used (an identity map). This field is ignored for services with clusterIP=None, and should be omitted or set equal to the 'port' field. More info: https://kubernetes.io/docs/concepts/services-networking/service/#defining-a-service

表 5-1754 io.k8s.api.core.v1.SessionAffinityConfig

参数	参数类型	描述
clientIP	io.k8s.api.core.v1.ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 5-1755 io.k8s.api.core.v1.ClientIPConfig

参数	参数类型	描述
timeoutSeconds	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 5-1756 io.k8s.api.core.v1.ServiceStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present.

表 5-1757 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-1758 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)

参数	参数类型	描述
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Service",
  "metadata": {
    "creationTimestamp": "2018-09-04T00:45:36Z",
    "labels": {
      "app": "redis"
    }
  },
  "name": "redis",
  "namespace": "namespace-test",
  "resourceVersion": "5146412",
  "selfLink": "/api/v1/namespaces/namespace-test/services/redis/status",
  "uid": "d6a1ce79-afdb-11e8-b6ef-f898ef6c78b4"
},
"spec": {
  "clusterIP": "10.247.212.210",
  "ports": [ {
    "name": "service0",
    "port": 8080,
    "protocol": "TCP",
    "targetPort": 80
  } ],
  "selector": {
    "app": "redis"
  },
  "sessionAffinity": "None",
  "type": "ClusterIP"
},
"status": {
  "loadBalancer": { }
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5 Deployment

5.5.1 查询用户所有 Deployments

功能介绍

查询用户所有Deployment。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/deployments

表 5-1759 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 5-1760 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1761 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.apps.v1.Deployment objects	Items is the list of Deployments.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata.

表 5-1762 io.k8s.api.apps.v1.Deployment

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-1763 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	参数类型	描述
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-1764 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-1765 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

参数	参数类型	描述
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-1766 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1767 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1768 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1769 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1770 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1771 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1772 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1773 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1774 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1775 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1776 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1777 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1778 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1779 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-1780 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1781 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-1782 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1783 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1784 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1785 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1786 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1787 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1788 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1789 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-1790 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-1791 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1792 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1793 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1794 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1795 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-1796 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1797 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1798 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1799 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-1800 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1801 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1802 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '..'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1803 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1804 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1805 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1806 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1807 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-1808 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1809 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-1810 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1811 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1812 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1813 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1814 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1815 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1816 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-1817 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1818 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1819 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1820 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1821 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1822 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1823 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1824 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1825 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1826 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1827 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1828 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1829 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1830 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1831 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1832 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1833 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1834 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1835 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1836 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1837 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1838 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1839 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1840 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1841 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1842 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1843 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1844 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1845 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1846 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1847 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1848 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1849 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1850 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1851 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1852 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1853 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1854 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1855 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1856 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1857 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1858 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1859 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1860 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1861 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1862 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-1863 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

表 5-1864 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "items": [ {
    "metadata": {
      "annotations": {
        "deployment.kubernetes.io/revision": "1",
        "service.potal.kubernetes.io/access-ip": "10.247.51.55:123",
        "service.potal.kubernetes.io/type": "ClusterIP"
      }
    }
  }
]
```

```
},
"creationTimestamp" : "2017-12-09T03:44:24Z",
"generation" : 2,
"labels" : {
  "app" : "nginx"
},
"name" : "nginx",
"namespace" : "default",
"resourceVersion" : "28579",
"selfLink" : "/apis/extensions/v1beta1/namespaces/default/deployments/nginx",
"uid" : "3f7846b5-dc93-11e7-9c19-fa163e2d897b"
},
"spec" : {
  "replicas" : 1,
  "selector" : {
    "matchLabels" : {
      "app" : "nginx"
    }
  },
  "strategy" : {
    "rollingUpdate" : {
      "maxSurge" : 1,
      "maxUnavailable" : 1
    },
    "type" : "RollingUpdate"
  },
  "template" : {
    "metadata" : {
      "creationTimestamp" : null,
      "labels" : {
        "app" : "nginx"
      }
    },
    "spec" : {
      "containers" : [ {
        "image" : "172.16.5.235:20202/test-01/mysql:v1",
        "imagePullPolicy" : "Always",
        "name" : "nginx",
        "resources" : { },
        "terminationMessagePath" : "/dev/termination-log",
        "terminationMessagePolicy" : "File"
      } ],
      "dnsPolicy" : "ClusterFirst",
      "restartPolicy" : "Always",
      "schedulerName" : "default-scheduler",
      "securityContext" : { }
    }
  },
  "status" : {
    "availableReplicas" : 1,
    "conditions" : [ {
      "lastTransitionTime" : "2017-12-09T03:44:24Z",
      "lastUpdateTime" : "2017-12-09T03:44:24Z",
      "message" : "Deployment has minimum availability.",
      "reason" : "MinimumReplicasAvailable",
      "status" : "True",
      "type" : "Available"
    } ],
    "observedGeneration" : 2,
    "readyReplicas" : 1,
    "replicas" : 1,
    "updatedReplicas" : 1
  }
}, {
  "metadata" : {
    "annotations" : {
      "deployment.kubernetes.io/revision" : "1"
    }
  },
}
```

```
"creationTimestamp" : "2017-12-13T03:13:22Z",
"generation" : 2,
"labels" : {
  "cce/appgroup" : "deploy-ex-test"
},
"name" : "deploy-ex-12130306",
"namespace" : "ns-12130306-s",
"resourceVersion" : "418771",
"selfLink" : "/apis/extensions/v1beta1/namespaces/ns-12130306-s/deployments/deploy-ex-12130306",
"uid" : "934db57d-dfb3-11e7-9c19-fa163e2d897b"
},
"spec" : {
  "replicas" : 1,
  "selector" : {
    "matchLabels" : {
      "cce/appgroup" : "deploy-ex-test"
    }
  },
  "strategy" : {
    "rollingUpdate" : {
      "maxSurge" : 1,
      "maxUnavailable" : 1
    },
    "type" : "RollingUpdate"
  },
  "template" : {
    "metadata" : {
      "creationTimestamp" : null,
      "labels" : {
        "cce/appgroup" : "deploy-ex-test"
      }
    },
    "spec" : {
      "containers" : [ {
        "image" : "172.16.5.235:20202/test/redis:latest",
        "imagePullPolicy" : "IfNotPresent",
        "name" : "deploycon-12130306",
        "resources" : { },
        "terminationMessagePath" : "/dev/termination-log",
        "terminationMessagePolicy" : "File"
      } ],
      "dnsPolicy" : "ClusterFirst",
      "restartPolicy" : "Always",
      "schedulerName" : "default-scheduler",
      "securityContext" : { }
    }
  },
  "status" : {
    "availableReplicas" : 1,
    "conditions" : [ {
      "lastTransitionTime" : "2017-12-13T03:13:22Z",
      "lastUpdateTime" : "2017-12-13T03:13:22Z",
      "message" : "Deployment has minimum availability.",
      "reason" : "MinimumReplicasAvailable",
      "status" : "True",
      "type" : "Available"
    } ],
    "observedGeneration" : 2,
    "readyReplicas" : 1,
    "replicas" : 1,
    "updatedReplicas" : 1
  }
},
"kind" : "DeploymentList",
"metadata" : {
  "resourceVersion" : "418820",
  "selfLink" : "/apis/extensions/v1beta1/deployments"
```

```
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.2 删除指定 namespace 下 Deployments

功能介绍

删除Namespace下所有Deployment。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/apps/v1/namespaces/{namespace}/deployments

表 5-1865 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1866 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1867 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-1868 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-1869 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-1870 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1871 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1872 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-1873 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

- 只删除Deployment（对应ReplicSet和Pod不删除）。

```
{
  "Kind": "DeleteOptions",
  "apiVersion": "v1",
  "propagationPolicy": "Orphan"
}
```

- 前台级联删除（按照Pod->ReplicaSet->Deployment的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Foreground"
}
```

- 后台级联删除（按照Deployment->ReplicaSet->Pod的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Background"
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "apps/v1",
  "items": null,
  "kind": "DeploymentList",
  "metadata": {
    "resourceVersion": "5039958",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/deployments"
  }
}
```


状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.3 查询指定 namespace 下的 Deployments

功能介绍

查询Namespace下所有Deployment的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/deployments

表 5-1874 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1875 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1876 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-1877 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.apps.v1.Deployment objects	Items is the list of Deployments.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata.

表 5-1878 io.k8s.api.apps.v1.Deployment

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-1879 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	参数类型	描述
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-1880 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-1881 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

参数	参数类型	描述
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-1882 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1883 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1884 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1885 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1886 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1887 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1888 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1889 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1890 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1891 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1892 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1893 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-1894 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-1895 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-1896 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1897 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-1898 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-1899 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-1900 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-1901 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-1902 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-1903 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1904 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-1905 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-1906 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-1907 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-1908 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-1909 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-1910 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1911 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-1912 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-1913 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-1914 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-1915 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-1916 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-1917 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-1918 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain ':'. Must not contain '*'. Must not contain '?'. Must not contain '['. Must not contain ']'. Must not contain '&'. Must not contain '='. Must not contain '". Must not contain ''.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-1919 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-1920 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-1921 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-1922 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1923 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-1924 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-1925 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-1926 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-1927 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-1928 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-1929 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-1930 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-1931 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-1932 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-1933 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-1934 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-1935 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1936 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-1937 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-1938 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-1939 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-1940 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-1941 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-1942 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-1943 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-1944 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-1945 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-1946 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-1947 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-1948 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-1949 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-1950 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-1951 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-1952 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-1953 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-1954 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-1955 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-1956 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-1957 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-1958 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-1959 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-1960 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-1961 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-1962 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-1963 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-1964 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-1965 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-1966 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-1967 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-1968 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-1969 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-1970 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-1971 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-1972 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-1973 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-1974 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-1975 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-1976 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-1977 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-1978 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-1979 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

表 5-1980 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "items": [ {
    "metadata": {
      "annotations": {
        "deployment.kubernetes.io/revision": "1"
      },
    },
    "creationTimestamp": "2018-09-03T12:58:07Z",
```

```
"generation" : 1,
"labels" : {
  "app" : "redis"
},
"name" : "deployment-test",
"namespace" : "namespace-test",
"resourceVersion" : "5036888",
"selfLink" : "/apis/apps/v1/namespaces/namespace-test/deployments/deployment-test",
"uid" : "010506c7-af79-11e8-b6ef-f898ef6c78b4"
},
"spec" : {
  "progressDeadlineSeconds" : 600,
  "replicas" : 1,
  "revisionHistoryLimit" : 2,
  "selector" : {
    "matchLabels" : {
      "app" : "redis"
    }
  },
  "strategy" : {
    "rollingUpdate" : {
      "maxSurge" : "25%",
      "maxUnavailable" : "25%"
    },
    "type" : "RollingUpdate"
  },
  "template" : {
    "metadata" : {
      "annotations" : {
        "cri.cci.io/container-type" : "secure-container"
      },
      "creationTimestamp" : null,
      "labels" : {
        "app" : "redis"
      }
    },
    "spec" : {
      "containers" : [ {
        "image" : "redis",
        "imagePullPolicy" : "IfNotPresent",
        "name" : "container-0",
        "resources" : {
          "limits" : {
            "cpu" : "500m",
            "memory" : "1Gi"
          },
          "requests" : {
            "cpu" : "500m",
            "memory" : "1Gi"
          }
        }
      } ],
      "terminationMessagePath" : "/dev/termination-log",
      "terminationMessagePolicy" : "File"
    },
    "dnsPolicy" : "ClusterFirst",
    "imagePullSecrets" : [ {
      "name" : "imagepull-secret"
    } ],
    "restartPolicy" : "Always",
    "schedulerName" : "default-scheduler",
    "securityContext" : { }
  }
},
"status" : {
  "availableReplicas" : 1,
  "conditions" : [ {
    "lastTransitionTime" : "2018-09-03T12:58:12Z",
    "lastUpdateTime" : "2018-09-03T12:58:12Z",
```

```

"message": "Deployment has minimum availability.",
"reason": "MinimumReplicasAvailable",
"status": "True",
"type": "Available"
}, {
"lastTransitionTime": "2018-09-03T12:58:07Z",
"lastUpdateTime": "2018-09-03T12:58:12Z",
"message": "ReplicaSet \"deployment-test-57f7cff77c\" has successfully\nprogressed.",
"reason": "NewReplicaSetAvailable",
"status": "True",
"type": "Progressing"
}],
"observedGeneration": 1,
"readyReplicas": 1,
"replicas": 1,
"updatedReplicas": 1
}
}],
"kind": "DeploymentList",
"metadata": {
"resourceVersion": "5038849",
"selfLink": "/apis/apps/v1/namespaces/namespace-test/deployments"
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.4 创建 Deployment

功能介绍

创建一个Deployment。

调用方法

请参见[如何调用API](#)。

URI

POST /apis/apps/v1/namespaces/{namespace}/deployments

表 5-1981 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-1982 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-1983 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-1984 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.

参数	是否必选	参数类型	描述
spec	否	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	否	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-1985 io.k8s.api.apps.v1.DeploymentSpec

参数	是否必选	参数类型	描述
minReadySeconds	否	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	否	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	否	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	否	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	否	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	是否必选	参数类型	描述
selector	是	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	否	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	是	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-1986 io.k8s.api.apps.v1.DeploymentStrategy

参数	是否必选	参数类型	描述
rollingUpdate	否	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	否	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-1987 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	是否必选	参数类型	描述
maxSurge	否	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	否	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-1988 io.k8s.api.core.v1.PodTemplateSpec

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-1989 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	是否必选	参数类型	描述
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	是否必选	参数类型	描述
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod

参数	是否必选	参数类型	描述
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	是否必选	参数类型	描述
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.

参数	是否必选	参数类型	描述
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy

参数	是否必选	参数类型	描述
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	否	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	是否必选	参数类型	描述
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	否	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br)". If not specified, the pod will not have a domainname at all.

参数	是否必选	参数类型	描述
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-1990 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	是否必选	参数类型	描述
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-1991 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-1992 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-1993 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-1994 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-1995 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-1996 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-1997 io.k8s.api.core.v1.PodAntiAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-1998 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-1999 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2000 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2001 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.
value	否	String	value is the value of the option

表 5-2002 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	是否必选	参数类型	描述
volumeDevices	否	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2003 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-2004 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2005 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2006 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2007 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2008 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2009 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-2010 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2011 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-2012 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-2013 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2014 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2015 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2016 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2017 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-2018 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2019 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2020 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-2021 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-2022 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-2023 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-2024 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain '!'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2025 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2026 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2027 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2028 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2029 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-2030 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2031 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-2032 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	否	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	否	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2033 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2034 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2035 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2036 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2037 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2038 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-2039 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2040 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2041 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2042 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2043 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2044 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2045 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2046 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2047 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2048 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2049 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2050 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2051 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-2052 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2053 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2054 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2055 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2056 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2057 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2058 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2059 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-2060 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2061 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2062 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface <code>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)</code> will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.

参数	是否必选	参数类型	描述
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2063 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) (Note that [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) may be empty, from the "" case in [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= 0 1 ... 9 [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br). ./topic/body/section/</pre>

参数	是否必选	参数类型	描述
			<p>table/tgroup/tbody/row/entry/p/br {""} (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "+" "-" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1</p>

参数	是否必选	参数类型	描述
			<p>in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-2064 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2065 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2066 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-2067 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-2068 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2069 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2070 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2071 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2072 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2073 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-2074 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-2075 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-2076 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2077 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2078 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2079 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2080 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2081 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2082 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2083 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmdk

表 5-2084 io.k8s.api.apps.v1.DeploymentStatus

参数	是否必选	参数类型	描述
availableReplicas	否	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.

参数	是否必选	参数类型	描述
collisionCount	否	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	否	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	否	Long	The generation observed by the deployment controller.
readyReplicas	否	Integer	Total number of ready pods targeted by this deployment.
replicas	否	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	否	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.
updatedReplicas	否	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2085 io.k8s.api.apps.v1.DeploymentCondition

参数	是否必选	参数类型	描述
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.
lastUpdateTime	否	String	The last time this condition was updated.

参数	是否必选	参数类型	描述
message	否	String	A human readable message indicating details about the transition.
reason	否	String	The reason for the condition's last transition.
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of deployment condition.

响应参数

状态码： 200

表 5-2086 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2087 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2088 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2089 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2090 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2091 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br].svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2092 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2093 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2094 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2095 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2096 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2097 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2098 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2099 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2100 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2101 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2102 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2103 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2104 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2105 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2106 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2107 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2108 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2109 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2110 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2111 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2112 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2113 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2114 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2115 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2116 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2117 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2118 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2119 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2120 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2121 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2122 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-2123 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-2124 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-2125 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-2126 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2127 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2128 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2129 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2130 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2131 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2132 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2133 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2134 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2135 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2136 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2137 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2138 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2139 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2140 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2141 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2142 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2143 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2144 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2145 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2146 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2147 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2148 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2149 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2150 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2151 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2152 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2153 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2154 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2155 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2156 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2157 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2158 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2159 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2160 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2161 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2162 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2163 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2164 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2165 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre> [/<topic ""="" "-"="" (br)="" (br).="" (br).)="" (br)[="" (international="" (note="" ...="" 1="" 9="" ::="Ki" <a="" <topic="" [="" be="" body="" br="" case="" ei="" empty,="" entry="" from="" gi="" href="http://physics.nist.gov/cuu/Units/binary.html" in="" may="" mi="" of="" p="" pi="" row="" section="" see:="" system="" table="" tbody="" tgroup="" that="" the="" ti="" units;="" {""}="" ="">http://physics.nist.gov/cuu/Units/binary.html) [/<topic <="" body="" pre="" section="" table="" tgroup=""> </topic></topic></pre>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-2166 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2167 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2168 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2169 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2170 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2171 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2172 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2173 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2174 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2175 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2176 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2177 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2178 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2179 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2180 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2181 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2182 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2183 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2184 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2185 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2186 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2187 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

状态码： 201

表 5-2188 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2189 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	参数类型	描述
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2190 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2191 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

参数	参数类型	描述
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2192 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2193 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2194 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2195 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2196 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2197 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2198 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2199 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2200 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2201 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2202 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2203 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2204 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2205 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2206 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2207 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2208 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2209 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2210 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2211 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2212 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2213 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2214 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2215 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2216 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2217 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2218 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2219 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2220 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2221 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2222 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2223 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2224 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-2225 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-2226 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-2227 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-2228 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '..'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2229 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2230 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2231 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2232 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2233 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2234 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2235 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2236 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2237 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2238 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2239 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2240 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2241 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2242 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2243 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2244 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2245 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2246 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2247 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2248 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2249 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2250 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2251 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2252 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2253 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2254 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2255 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2256 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2257 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2258 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2259 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2260 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2261 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2262 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2263 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2264 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2265 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2266 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2267 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre> [/<topic ""="" "-"="" (br)="" (br).="" (br).)="" (br).[="" (br)[="" (international="" (note="" .="" ...="" 1="" 9="" ::="Ki" <a="" <topic="" [="" be="" body="" br="" case="" ei="" empty,="" entry="" from="" gi="" href="http://physics.nist.gov/cuu/Units/binary.html" in="" may="" mi="" of="" p="" pi="" row="" section="" see:="" system="" table="" tbody="" tgroup="" that="" the="" ti="" units;="" {""}="" ="">http://physics.nist.gov/cuu/Units/binary.html) [/<topic <="" body="" pre="" section="" table="" tgroup=""> </topic></topic></pre>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-2268 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2269 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2270 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2271 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2272 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2273 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2274 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2275 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2276 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2277 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2278 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2279 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2280 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2281 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2282 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2283 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2284 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2285 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2286 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2287 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2288 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2289 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

状态码： 202

表 5-2290 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2291 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	参数类型	描述
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2292 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2293 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

参数	参数类型	描述
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2294 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2295 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2296 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2297 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2298 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2299 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2300 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2301 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2302 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2303 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2304 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2305 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2306 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2307 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2308 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2309 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2310 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2311 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2312 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2313 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2314 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2315 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2316 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2317 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2318 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2319 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2320 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2321 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2322 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2323 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2324 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2325 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2326 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2331 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2332 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2333 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2334 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2335 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2336 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2337 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2338 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entryptpoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2339 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2340 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2341 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2342 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2343 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2344 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2345 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2346 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2347 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2348 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2349 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2350 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2351 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2352 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2353 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2354 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2355 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2356 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2357 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2358 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2359 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2360 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2361 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2362 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2363 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2364 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2365 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2366 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2367 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2368 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2369 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre> [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br)[topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) (Note that [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) may be empty, from the "" case in [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br).) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= 0 1 ... 9 [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br)[topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br). . [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= "+" "-" [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br)[topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) [topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [topic/body/section/table/tgroup/ </pre>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-2370 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2371 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2372 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2373 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2374 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2375 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2376 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2377 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2378 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2379 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2380 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2381 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2382 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2383 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2384 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2385 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2386 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2387 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2388 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2389 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2390 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2391 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

请求示例

- 创建一个名为"deployment-test"的普通Deployment，使用redis镜像创建一个Pod，每个Pod占用500m core CPU、1024Mi内存。

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "name": "deployment-test"
  },
  "spec": {
    "replicas": 1,
    "selector": {
      "matchLabels": {
        "app": "redis"
      }
    },
    "template": {
      "metadata": {
        "labels": {
          "app": "redis"
        }
      },
      "spec": {
        "containers": [ {
          "image": "redis",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1024Mi"
            }
          }
        }
      ]
    }
  }
}
```



```
    "requests": {
      "cpu": "500m",
      "memory": "1024Mi"
    }
  },
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "priority": 0
}
}
```

- 创建使用GPU的Deployment，NVIDIA GPU的驱动版本为418.126。

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "name": "test-app",
    "namespace": "test-namespace"
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "matchLabels": {
        "app": "test-app"
      }
    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/gpu-driver": "gpu-418.126"
        },
        "labels": {
          "app": "test-app"
        }
      },
      "spec": {
        "containers": [ {
          "image": "library/nginx:stable-alpine-perl",
          "lifecycle": { },
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "4",
              "memory": "32Gi",
              "nvidia.com/gpu-tesla-v100-16GB": "1"
            },
            "requests": {
              "cpu": "4",
              "memory": "32Gi",
              "nvidia.com/gpu-tesla-v100-16GB": "1"
            }
          }
        }
      ],
      "imagePullSecrets": [ {
        "name": "imagepull-secret"
      } ],
      "restartPolicy": "Always"
    }
  }
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "creationTimestamp": "2018-09-06T03:39:32Z",
    "generation": 1,
    "labels": {
      "app": "redis"
    },
    "name": "deployment-test",
    "namespace": "namespace-test",
    "resourceVersion": "5630832",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/deployments/deployment-test",
    "uid": "777dce52-b186-11e8-8cb0-c81fbe371a17"
  },
  "spec": {
    "progressDeadlineSeconds": 600,
    "replicas": 1,
    "revisionHistoryLimit": 10,
    "selector": {
      "matchLabels": {
        "app": "redis"
      }
    },
    "strategy": {
      "rollingUpdate": {
        "maxSurge": "25%",
        "maxUnavailable": "25%"
      },
      "type": "RollingUpdate"
    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "app": "redis"
        }
      },
      "spec": {
        "containers": [ {
          "image": "redis",
          "imagePullPolicy": "IfNotPresent",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1Gi"
            },
            "requests": {
              "cpu": "500m",
              "memory": "1Gi"
            }
          }
        } ],
        "terminationMessagePath": "/dev/termination-log",
        "terminationMessagePolicy": "File"
      },
      "dnsPolicy": "ClusterFirst",
      "imagePullSecrets": [ {
        "name": "imagepull-secret"
      } ],
      "restartPolicy": "Always",
      "schedulerName": "default-scheduler",
      "securityContext": { }
    }
  }
}
```

```
},  
"status": {}  
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.5 删除 Deployment

功能介绍

删除Deployment。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/apps/v1/namespaces/{namespace}/deployments/{name}

表 5-2392 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Deployment
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2393 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2394 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-2395 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-2396 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-2397 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2398 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2399 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-2400 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-2401 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2402 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2403 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-2404 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

- 只删除Deployment（对应ReplicSet和Pod不删除）。

```
{
  "Kind": "DeleteOptions",
  "apiVersion": "v1",
  "propagationPolicy": "Orphan"
}
```

- 前台级联删除（按照Pod->ReplicaSet->Deployment的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Foreground"
}
```

- 后台级联删除（按照Deployment->ReplicaSet->Pod的顺序进行删除）

```
{
  "apiVersion": "v1",
```

```
"kind": "DeleteOptions",  
"propagationPolicy": "Background"  
}
```

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "v1",  
  "code": 200,  
  "details": {  
    "group": "extensions",  
    "kind": "deployments",  
    "name": "deploy-12130306",  
    "uid": "27072a31-dfb3-11e7-9c19-fa163e2d897b"  
  },  
  "kind": "Status",  
  "metadata": { },  
  "status": "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.6 查询 Deployment

功能介绍

查询Deployment的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}

表 5-2405 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Deployment
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2406 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2407 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-2408 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2409 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2410 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2411 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2412 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2413 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2414 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2415 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2416 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2417 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2418 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2419 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2420 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2421 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2422 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2423 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2424 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2425 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2426 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2427 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2428 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2429 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2430 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2431 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2432 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2433 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2434 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2435 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2436 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2437 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2438 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2439 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2440 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2441 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2442 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2443 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2444 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-2445 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-2446 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-2447 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-2448 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2449 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2450 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2451 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2452 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2453 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2454 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2455 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2456 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2457 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2458 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2459 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2460 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2461 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2462 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2463 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2464 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2465 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2466 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2467 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2468 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2469 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2470 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2471 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2472 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2473 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2474 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2475 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2476 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2477 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2478 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2479 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2480 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2481 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2482 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2483 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2484 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2485 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2486 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2487 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-2488 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2489 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2490 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2491 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2492 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2493 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2494 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2495 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2496 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2497 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2498 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2499 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2500 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2501 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2502 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2503 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2504 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2505 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2506 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2507 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2508 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2509 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "annotations": {
      "deployment.kubernetes.io/revision": "1"
    },
    "creationTimestamp": "2018-09-03T12:58:07Z",
    "generation": 1,
    "labels": {
      "app": "redis"
    },
    "name": "deployment-test",
    "namespace": "namespace-test",
    "resourceVersion": "5036888",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/deployments/deployment-test",
    "uid": "010506c7-af79-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "progressDeadlineSeconds": 600,
    "replicas": 1,

```

```
"revisionHistoryLimit" : 2,
"selector" : {
  "matchLabels" : {
    "app" : "redis"
  }
},
"strategy" : {
  "rollingUpdate" : {
    "maxSurge" : "25%",
    "maxUnavailable" : "25%"
  },
  "type" : "RollingUpdate"
},
"template" : {
  "metadata" : {
    "annotations" : {
      "cri.cci.io/container-type" : "secure-container"
    },
    "creationTimestamp" : null,
    "labels" : {
      "app" : "redis"
    }
  },
  "spec" : {
    "containers" : [ {
      "image" : "redis",
      "imagePullPolicy" : "IfNotPresent",
      "name" : "container-0",
      "resources" : {
        "limits" : {
          "cpu" : "500m",
          "memory" : "1Gi"
        },
        "requests" : {
          "cpu" : "500m",
          "memory" : "1Gi"
        }
      },
      "terminationMessagePath" : "/dev/termination-log",
      "terminationMessagePolicy" : "File"
    } ],
    "dnsPolicy" : "ClusterFirst",
    "imagePullSecrets" : [ {
      "name" : "imagepull-secret"
    } ],
    "restartPolicy" : "Always",
    "schedulerName" : "default-scheduler",
    "securityContext" : { }
  }
},
"status" : {
  "availableReplicas" : 1,
  "conditions" : [ {
    "lastTransitionTime" : "2018-09-03T12:58:12Z",
    "lastUpdateTime" : "2018-09-03T12:58:12Z",
    "message" : "Deployment has minimum availability.",
    "reason" : "MinimumReplicasAvailable",
    "status" : "True",
    "type" : "Available"
  }, {
    "lastTransitionTime" : "2018-09-03T12:58:07Z",
    "lastUpdateTime" : "2018-09-03T12:58:12Z",
    "message" : "ReplicaSet \"deployment-test-57f7cff77c\" has successfully\nprogressed.",
    "reason" : "NewReplicaSetAvailable",
    "status" : "True",
    "type" : "Progressing"
  } ],
  "observedGeneration" : 1,
```



```
"readyReplicas" : 1,  
"replicas" : 1,  
"updatedReplicas" : 1  
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.7 更新 Deployment

功能介绍

更新Deployment。

The following fields can be updated:

- metadata.labels
- metadata.generateName
- metadata.annotations
- spec.template
- spec.replicas
- spec.revisionHistoryLimit
- spec.progressDeadlineSeconds

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/apps/v1/namespaces/{namespace}/deployments/{name}

表 5-2510 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Deployment
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2511 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

参数	是否必选	参数类型	描述
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2512 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-2513 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-2514 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2515 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.

参数	参数类型	描述
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2516 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2517 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2518 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2519 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.

参数	参数类型	描述
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	参数类型	描述
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

参数	参数类型	描述
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://github.com/kubernetes/enhancements/blob/master/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccountName	String	DeprecatedServiceAccountName is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	参数类型	描述
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2520 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2521 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2522 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2523 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2524 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2525 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2526 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2527 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2528 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2529 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2530 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2531 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2532 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2533 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2534 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2535 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2536 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2537 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2538 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2539 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2540 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2541 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2542 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2543 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2544 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2545 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2546 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2547 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2548 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2549 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2550 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2555 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2556 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2557 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2558 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2559 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2560 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2561 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2562 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2563 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2564 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2565 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2566 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2567 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2568 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2569 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2570 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2571 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2572 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2573 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2574 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2575 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2576 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2577 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2578 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2579 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2580 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2581 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2582 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2583 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2584 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2585 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2586 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2587 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2588 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2589 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2590 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2591 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2592 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2593 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-2594 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2595 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2596 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2597 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2598 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2599 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2600 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2601 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2602 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2603 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2604 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2605 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2606 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2607 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2608 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2609 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2610 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2611 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2612 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2613 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2614 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2615 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

请求示例

更新Deployment中的labels值。

```
{
  "metadata": {
    "labels": {
      "app": "deployment-test2"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "annotations": {
      "deployment.kubernetes.io/revision": "2"
    },
    "creationTimestamp": "2018-09-06T03:39:32Z",
    "generation": 2,
    "labels": {
      "app": "deployment-test2"
    },
    "name": "deployment-test",
```

```
"namespace" : "namespace-test",
"resourceVersion" : "5657176",
"selfLink" : "/apis/apps/v1/namespaces/namespace-test/deployments/deployment-test",
"uid" : "777dce52-b186-11e8-8cb0-c81fbe371a17"
},
"spec" : {
  "progressDeadlineSeconds" : 600,
  "replicas" : 2,
  "revisionHistoryLimit" : 10,
  "selector" : {
    "matchLabels" : {
      "app" : "redis"
    }
  },
  "strategy" : {
    "rollingUpdate" : {
      "maxSurge" : "25%",
      "maxUnavailable" : "25%"
    },
    "type" : "RollingUpdate"
  },
  "template" : {
    "metadata" : {
      "creationTimestamp" : null,
      "labels" : {
        "app" : "redis"
      }
    },
    "spec" : {
      "containers" : [ {
        "image" : "k8s.gcr.io:20202/cci/redis:v1",
        "imagePullPolicy" : "IfNotPresent",
        "name" : "container-0",
        "resources" : {
          "limits" : {
            "cpu" : "500m",
            "memory" : "1Gi"
          },
          "requests" : {
            "cpu" : "500m",
            "memory" : "1Gi"
          }
        },
        "terminationMessagePath" : "/dev/termination-log",
        "terminationMessagePolicy" : "File"
      } ],
      "dnsPolicy" : "ClusterFirst",
      "imagePullSecrets" : [ {
        "name" : "imagepull-secret"
      } ],
      "restartPolicy" : "Always",
      "schedulerName" : "default-scheduler",
      "securityContext" : { }
    }
  },
  "status" : {
    "availableReplicas" : 2,
    "conditions" : [ {
      "lastTransitionTime" : "2018-09-06T04:14:14Z",
      "lastUpdateTime" : "2018-09-06T04:14:14Z",
      "message" : "Deployment has minimum availability.",
      "reason" : "MinimumReplicasAvailable",
      "status" : "True",
      "type" : "Available"
    }, {
      "lastTransitionTime" : "2018-09-06T03:39:32Z",
      "lastUpdateTime" : "2018-09-06T04:14:24Z",
      "message" : "ReplicaSet \"deployment-test-68585dfddb\" has successfully\nprogressed."
    } ]
  }
}
```

```
"reason" : "NewReplicaSetAvailable",
"status" : "True",
"type" : "Progressing"
}],
"observedGeneration" : 2,
"readyReplicas" : 2,
"replicas" : 2,
"updatedReplicas" : 2
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.8 替换 Deployment

功能介绍

替换Deployment。

其中以下字段支持更新：

- metadata.labels
- metadata.generateName
- metadata.annotations
- spec.template
- spec.replicas

- spec.revisionHistoryLimit
- spec.progressDeadlineSeconds

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}

表 5-2616 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Deployment
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2617 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2618 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-2619 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.

参数	是否必选	参数类型	描述
spec	否	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	否	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2620 io.k8s.api.apps.v1.DeploymentSpec

参数	是否必选	参数类型	描述
minReadySeconds	否	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	否	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	否	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	否	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	否	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	是否必选	参数类型	描述
selector	是	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	否	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	是	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2621 io.k8s.api.apps.v1.DeploymentStrategy

参数	是否必选	参数类型	描述
rollingUpdate	否	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	否	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2622 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	是否必选	参数类型	描述
maxSurge	否	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	否	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2623 io.k8s.api.core.v1.PodTemplateSpec

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2624 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	是否必选	参数类型	描述
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	是否必选	参数类型	描述
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod

参数	是否必选	参数类型	描述
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	是否必选	参数类型	描述
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.

参数	是否必选	参数类型	描述
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy

参数	是否必选	参数类型	描述
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	否	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	是否必选	参数类型	描述
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	否	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.

参数	是否必选	参数类型	描述
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2625 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	是否必选	参数类型	描述
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2626 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2627 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2628 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2629 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2630 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2631 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-2632 io.k8s.api.core.v1.PodAntiAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2633 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2634 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2635 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2636 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.
value	否	String	value is the value of the option

表 5-2637 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	是否必选	参数类型	描述
volumeDevices	否	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2638 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-2639 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2640 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2641 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2642 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2643 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2644 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-2645 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2646 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-2647 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-2648 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2649 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2650 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2651 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2652 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-2653 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2654 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2655 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-2656 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-2657 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-2658 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-2659 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain '!'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2660 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2661 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2662 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2663 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2664 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-2665 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2666 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-2667 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
gmsaCredentialSpecName	否	String	<code>GMSACredentialSpecName</code> is the name of the GMSA credential spec to use.
runAsUserName	否	String	The <code>UserName</code> in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in <code>PodSecurityContext</code> . If set in both <code>SecurityContext</code> and <code>PodSecurityContext</code> , the value specified in <code>SecurityContext</code> takes precedence.

表 5-2668 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are <code>NoSchedule</code> , <code>PreferNoSchedule</code> and <code>NoExecute</code> .
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be <code>Exists</code> ; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2669 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2670 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2671 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2672 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2673 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-2674 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2675 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2676 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2677 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2678 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2679 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2680 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2681 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2682 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2683 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2684 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2685 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2686 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-2687 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2688 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2689 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2690 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2691 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2692 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2693 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2694 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-2695 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2696 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2697 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.

参数	是否必选	参数类型	描述
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2698 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or</p>

参数	是否必选	参数类型	描述
			<p>suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-2699 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2700 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2701 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-2702 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-2703 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2704 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2705 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2706 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2707 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2708 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-2709 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-2710 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-2711 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2712 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2713 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2714 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2715 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2716 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2717 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2718 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmdk

表 5-2719 io.k8s.api.apps.v1.DeploymentStatus

参数	是否必选	参数类型	描述
availableReplicas	否	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.

参数	是否必选	参数类型	描述
collisionCount	否	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	否	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	否	Long	The generation observed by the deployment controller.
readyReplicas	否	Integer	Total number of ready pods targeted by this deployment.
replicas	否	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	否	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.
updatedReplicas	否	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2720 io.k8s.api.apps.v1.DeploymentCondition

参数	是否必选	参数类型	描述
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.
lastUpdateTime	否	String	The last time this condition was updated.

参数	是否必选	参数类型	描述
message	否	String	A human readable message indicating details about the transition.
reason	否	String	The reason for the condition's last transition.
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of deployment condition.

响应参数

状态码： 200

表 5-2721 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2722 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2723 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2724 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2725 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2726 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2727 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2728 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2729 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2730 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2731 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2732 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2733 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2734 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2735 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2736 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2737 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2738 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2739 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2740 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2741 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2742 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2743 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2744 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2745 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2746 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2747 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2748 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2749 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2750 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2751 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2752 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2753 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2754 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2755 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2756 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2757 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2762 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2763 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2764 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2765 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2766 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2767 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2768 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2769 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2770 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2771 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2772 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2773 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2774 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2775 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2776 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2777 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2778 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2779 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2780 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2781 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2782 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2783 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2784 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2785 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2786 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2787 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2788 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2789 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2790 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2791 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2792 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2793 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2794 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2795 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2796 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2797 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2798 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2799 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2800 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-2801 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2802 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2803 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2804 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2805 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2806 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2807 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2808 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2809 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2810 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2811 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2812 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2813 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2814 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2815 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2816 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2817 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2818 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2819 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2820 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2821 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2822 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

状态码： 201

表 5-2823 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2824 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.

参数	参数类型	描述
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2825 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2826 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

参数	参数类型	描述
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2827 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2828 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2829 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2830 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2831 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2832 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2833 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2834 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2835 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2836 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2837 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2838 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2839 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2840 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2841 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2842 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2843 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2844 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2845 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2846 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2847 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2848 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2849 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2850 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2851 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2852 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2853 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2854 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-2855 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2856 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-2857 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-2858 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-2859 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-2864 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-2865 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-2866 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-2867 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2868 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-2869 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-2870 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-2871 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-2872 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-2873 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-2874 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-2875 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-2876 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-2877 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-2878 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-2879 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-2880 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2881 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-2882 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-2883 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-2884 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-2885 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-2886 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2887 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2888 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2889 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-2890 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-2891 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-2892 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-2893 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-2894 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-2895 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-2896 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-2897 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-2898 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-2899 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-2900 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-2901 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-2902 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-2903 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-2904 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-2905 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-2906 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-2907 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-2908 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-2909 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-2910 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-2911 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-2912 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-2913 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-2914 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2915 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-2916 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-2917 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-2918 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-2919 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-2920 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-2921 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-2922 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-2923 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-2924 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

请求示例

将已创建Deployment中的镜像替换为 “ ...:20202/cci/redis:V1 ”。

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "name": "deployment-test"
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "matchLabels": {
        "app": "redis"
      }
    },
    "template": {
      "metadata": {
        "labels": {
          "app": "redis"
        }
      },
      "spec": {
        "containers": [ {
          "image": "***:20202/cci/redis:V1",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1024Mi"
            },
            "requests": {
```

```
        "cpu": "500m",
        "memory": "1024Mi"
      }
    }
  },
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ]
}
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "creationTimestamp": "2018-09-06T03:39:32Z",
    "generation": 2,
    "labels": {
      "app": "redis"
    },
    "name": "deployment-test",
    "namespace": "namespace-test",
    "resourceVersion": "5636210",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/deployments/deployment-test",
    "uid": "777dce52-b186-11e8-8cb0-c81f371a17"
  },
  "spec": {
    "progressDeadlineSeconds": 600,
    "replicas": 2,
    "revisionHistoryLimit": 10,
    "selector": {
      "matchLabels": {
        "app": "redis"
      }
    },
    "strategy": {
      "rollingUpdate": {
        "maxSurge": "25%",
        "maxUnavailable": "25%"
      },
      "type": "RollingUpdate"
    },
    "template": {
      "metadata": {
        "creationTimestamp": null,
        "labels": {
          "app": "redis"
        }
      },
      "spec": {
        "containers": [ {
          "image": "redis",
          "imagePullPolicy": "IfNotPresent",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1Gi"
            },
            "requests": {
              "cpu": "500m",
```

```
    "memory": "1Gi"
  }
},
"terminationMessagePath": "/dev/termination-log",
"terminationMessagePolicy": "File"
}],
"dnsPolicy": "ClusterFirst",
"imagePullSecrets": [ {
  "name": "imagepull-secret"
}],
"restartPolicy": "Always",
"schedulerName": "default-scheduler",
"securityContext": { }
}
},
"status": {
  "availableReplicas": 1,
  "conditions": [ {
    "lastTransitionTime": "2018-09-06T03:39:40Z",
    "lastUpdateTime": "2018-09-06T03:39:40Z",
    "message": "Deployment has minimum availability.",
    "reason": "MinimumReplicasAvailable",
    "status": "True",
    "type": "Available"
  }, {
    "lastTransitionTime": "2018-09-06T03:39:32Z",
    "lastUpdateTime": "2018-09-06T03:39:40Z",
    "message": "ReplicaSet \"deployment-test-865578b586\" has successfully\nprogressed.",
    "reason": "NewReplicaSetAvailable",
    "status": "True",
    "type": "Progressing"
  } ],
  "observedGeneration": 1,
  "readyReplicas": 1,
  "replicas": 1,
  "updatedReplicas": 1
}
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid

状态码	描述
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.9 查询 Deployment 的伸缩操作

功能介绍

查询Deployment的伸缩操作

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale

表 5-2925 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Scale
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2926 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2927 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-2928 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata; More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata .
spec	io.k8s.api.autoscaling.v1.ScaleSpec object	defines the behavior of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status .
status	io.k8s.api.autoscaling.v1.ScaleStatus object	current status of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status . Read-only.

表 5-2929 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2930 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2931 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2932 io.k8s.api.autoscaling.v1.ScaleSpec

参数	参数类型	描述
replicas	Integer	desired number of instances for the scaled object.

表 5-2933 io.k8s.api.autoscaling.v1.ScaleStatus

参数	参数类型	描述
replicas	Integer	actual number of observed instances of the scaled object.
selector	String	label query over pods that should match the replicas count. This is same as the label selector but in the string format to avoid introspection by clients. The string will be in the same format as the query-param syntax. More info about label selectors: http://kubernetes.io/docs/user-guide/labels#label-selectors

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "autoscaling/v1",
  "kind": "Scale",
  "metadata": {
    "creationTimestamp": "2022-08-30T09:23:43Z",
    "name": "nginx",
    "namespace": "test-api",
    "resourceVersion": "34017936",
    "selfLink": "/apis/apps/v1/namespaces/test-api/deployments/nginx/scale",
    "uid": "3bd1306c-ac3e-41c2-81ab-14c79fa006d3"
  },
  "spec": {
    "replicas": 2
  },
  "status": {
    "replicas": 2,
    "selector": "app=nginx"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.10 更新 Deployment 的伸缩操作

功能介绍

This API is used to partially update scale of the specified Scale.

The following fields can be updated:

- metadata.resourceVersion
- metadata.creationTimestamp
- spec.replicas

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale

表 5-2934 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Scale
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2935 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2936 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-2937 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-2938 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata; More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata .
spec	io.k8s.api.autoscaling.v1.ScaleSpec object	defines the behavior of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status .
status	io.k8s.api.autoscaling.v1.ScaleStatus object	current status of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status . Read-only.

表 5-2939 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2940 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2941 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2942 io.k8s.api.autoscaling.v1.ScaleSpec

参数	参数类型	描述
replicas	Integer	desired number of instances for the scaled object.

表 5-2943 io.k8s.api.autoscaling.v1.ScaleStatus

参数	参数类型	描述
replicas	Integer	actual number of observed instances of the scaled object.
selector	String	label query over pods that should match the replicas count. This is same as the label selector but in the string format to avoid introspection by clients. The string will be in the same format as the query-param syntax. More info about label selectors: http://kubernetes.io/docs/user-guide/labels#label-selectors

请求示例

更新Deployment的伸缩操作。

```
[{
  "op": "add",
  "path": "/spec/replicas",
  "value": 2
}]
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "autoscaling/v1",
  "kind": "Scale",
  "metadata": {
    "creationTimestamp": "2022-08-30T09:23:43Z",
    "name": "nginx",
    "namespace": "test-api",
    "resourceVersion": "34017936",
    "selfLink": "/apis/apps/v1/namespaces/test-api/deployments/nginx/scale",
    "uid": "3bd1306c-ac3e-41c2-81ab-14c79fa006d3"
  },
  "spec": {
    "replicas": 2
  },
  "status": {
    "replicas": 2,
    "selector": "app=nginx"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest

状态码	描述
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.11 替换 Deployment 的伸缩操作

功能介绍

This API is used to replace scale of the specified Scale.

The following fields can be updated:

- metadata.resourceVersion
- metadata.creationTimestamp
- spec.replicas

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale

表 5-2944 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Scale

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2945 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2946 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-2947 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata; More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata .
spec	否	io.k8s.api.autoscaling.v1.ScaleSpec object	defines the behavior of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status .
status	否	io.k8s.api.autoscaling.v1.ScaleStatus object	current status of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status . Read-only.

表 5-2948 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.
labels	否	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2949 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2950 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2951 io.k8s.api.autoscaling.v1.ScaleSpec

参数	是否必选	参数类型	描述
replicas	否	Integer	desired number of instances for the scaled object.

表 5-2952 io.k8s.api.autoscaling.v1.ScaleStatus

参数	是否必选	参数类型	描述
replicas	是	Integer	actual number of observed instances of the scaled object.

参数	是否必选	参数类型	描述
selector	否	String	label query over pods that should match the replicas count. This is same as the label selector but in the string format to avoid introspection by clients. The string will be in the same format as the query-param syntax. More info about label selectors: http://kubernetes.io/docs/user-guide/labels#label-selectors

响应参数

状态码： 200

表 5-2953 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata; More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata .
spec	io.k8s.api.autoscaling.v1.ScaleSpec object	defines the behavior of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status .

参数	参数类型	描述
status	io.k8s.api.autoscaling.v1.ScaleStatus object	current status of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status . Read-only.

表 5-2954 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-2955 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2956 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2957 io.k8s.api.autoscaling.v1.ScaleSpec

参数	参数类型	描述
replicas	Integer	desired number of instances for the scaled object.

表 5-2958 io.k8s.api.autoscaling.v1.ScaleStatus

参数	参数类型	描述
replicas	Integer	actual number of observed instances of the scaled object.
selector	String	label query over pods that should match the replicas count. This is same as the label selector but in the string format to avoid introspection by clients. The string will be in the same format as the query-param syntax. More info about label selectors: http://kubernetes.io/docs/user-guide/labels#label-selectors

状态码: 201

表 5-2959 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata; More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata .
spec	io.k8s.api.autoscaling.v1.ScaleSpec object	defines the behavior of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status .
status	io.k8s.api.autoscaling.v1.ScaleStatus object	current status of the scale. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status . Read-only.

表 5-2960 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2961 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-2962 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-2963 io.k8s.api.autoscaling.v1.ScaleSpec

参数	参数类型	描述
replicas	Integer	desired number of instances for the scaled object.

表 5-2964 io.k8s.api.autoscaling.v1.ScaleStatus

参数	参数类型	描述
replicas	Integer	actual number of observed instances of the scaled object.
selector	String	label query over pods that should match the replicas count. This is same as the label selector but in the string format to avoid introspection by clients. The string will be in the same format as the query-param syntax. More info about label selectors: http://kubernetes.io/docs/user-guide/labels#label-selectors

请求示例

替换Deployment的伸缩操作。

```
{
  "apiVersion": "autoscaling/v1",
  "kind": "Scale",
  "metadata": {
    "creationTimestamp": "2022-08-30T09:23:43Z",
    "name": "nginx",
    "namespace": "test-api",
    "resourceVersion": "34017936",
    "selfLink": "/apis/apps/v1/namespaces/test-api/deployments/nginx/scale",
    "uid": "3bd1306c-ac3e-41c2-81ab-14c79fa006d3"
  },
  "spec": {
    "replicas": 2
  },
  "status": {
    "replicas": 2,
    "selector": "app=nginx"
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "autoscaling/v1",
  "kind": "Scale",
  "metadata": {
    "creationTimestamp": "2022-08-30T09:23:43Z",
    "name": "nginx",
    "namespace": "test-api",
    "resourceVersion": "34017936",
    "selfLink": "/apis/apps/v1/namespaces/test-api/deployments/nginx/scale",
    "uid": "3bd1306c-ac3e-41c2-81ab-14c79fa006d3"
  },
  "spec": {
    "replicas": 2
  },
  "status": {
    "replicas": 2,
  }
}
```

```
"selector" : "app=nginx"  
}  
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.5.12 查询 Deployment 状态

功能介绍

查询Deployment的状态。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status

表 5-2965 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Deployment
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-2966 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-2967 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-2968 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object metadata.
spec	io.k8s.api.apps.v1.DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	io.k8s.api.apps.v1.DeploymentStatus object	Most recently observed status of the Deployment.

表 5-2969 io.k8s.api.apps.v1.DeploymentSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
paused	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.

参数	参数类型	描述
revisionHistoryLimit	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment. It must match the pod template's labels.
strategy	io.k8s.api.apps.v1.DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	io.k8s.api.core.v1.PodTemplateSpec object	Template describes the pods that will be created.

表 5-2970 io.k8s.api.apps.v1.DeploymentStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 5-2971 io.k8s.api.apps.v1.RollingUpdateDeployment

参数	参数类型	描述
maxSurge	String	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This can not be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new ReplicaSet can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new ReplicaSet can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.
maxUnavailable	String	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This can not be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old ReplicaSet can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old ReplicaSet can be scaled down further, followed by scaling up the new ReplicaSet, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 5-2972 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-2973 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.

参数	参数类型	描述
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	参数类型	描述
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

参数	参数类型	描述
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	参数类型	描述
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-2974 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-2975 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-2976 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-2977 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-2978 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-2979 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-2980 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2981 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-2982 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-2983 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-2984 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-2985 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-2986 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2987 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-2988 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-2989 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-2990 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-2991 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-2992 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-2993 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-2994 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-2995 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-2996 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-2997 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-2998 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-2999 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-3000 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-3001 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-3002 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-3003 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-3004 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-3009 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-3010 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-3011 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-3012 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-3013 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-3014 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-3015 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-3016 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-3017 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-3018 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-3019 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-3020 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-3021 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-3022 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-3023 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-3024 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-3025 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-3026 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-3027 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-3028 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-3029 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-3030 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-3031 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3032 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3033 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3034 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3035 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3036 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3037 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3038 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3039 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-3040 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-3041 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-3042 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-3043 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-3044 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-3045 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-3046 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-3047 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-3048 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-3049 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-3050 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-3051 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-3052 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-3053 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-3054 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-3055 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-3056 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-3057 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-3058 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-3059 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-3060 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-3061 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-3062 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-3063 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-3064 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-3065 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-3066 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-3067 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-3068 io.k8s.api.apps.v1.DeploymentStatus

参数	参数类型	描述
availableReplicas	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	Array of io.k8s.api.apps.v1.DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	Long	The generation observed by the deployment controller.
readyReplicas	Integer	Total number of ready pods targeted by this deployment.
replicas	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	Integer	Total number of unavailable pods targeted by this deployment. This is the total number of pods that are still required for the deployment to have 100% available capacity. They may either be pods that are running but not yet available or pods that still have not been created.

参数	参数类型	描述
updatedReplicas	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 5-3069 io.k8s.api.apps.v1.DeploymentCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
lastUpdateTime	String	The last time this condition was updated.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of deployment condition.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "Deployment",
  "metadata": {
    "annotations": {
      "deployment.kubernetes.io/revision": "1"
    },
    "creationTimestamp": "2018-09-06T03:39:32Z",
    "generation": 1,
    "labels": {
      "app": "redis"
    },
    "name": "deployment-test",
    "namespace": "namespace-test",
    "resourceVersion": "5630865",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/deployments/deployment-test/status",
    "uid": "777dce52-b186-11e8-8cb0-c81fbe371a17"
  },
  "spec": {
    "progressDeadlineSeconds": 600,
    "replicas": 1,

```

```
"revisionHistoryLimit" : 10,
"selector" : {
  "matchLabels" : {
    "app" : "redis"
  }
},
"strategy" : {
  "rollingUpdate" : {
    "maxSurge" : "25%",
    "maxUnavailable" : "25%"
  },
  "type" : "RollingUpdate"
},
"template" : {
  "metadata" : {
    "annotations" : {
      "cri.cci.io/container-type" : "secure-container"
    },
    "creationTimestamp" : null,
    "labels" : {
      "app" : "redis"
    }
  },
  "spec" : {
    "containers" : [ {
      "image" : "redis",
      "imagePullPolicy" : "IfNotPresent",
      "name" : "container-0",
      "resources" : {
        "limits" : {
          "cpu" : "500m",
          "memory" : "1Gi"
        },
        "requests" : {
          "cpu" : "500m",
          "memory" : "1Gi"
        }
      },
      "terminationMessagePath" : "/dev/termination-log",
      "terminationMessagePolicy" : "File"
    } ],
    "dnsPolicy" : "ClusterFirst",
    "imagePullSecrets" : [ {
      "name" : "imagepull-secret"
    } ],
    "restartPolicy" : "Always",
    "schedulerName" : "default-scheduler",
    "securityContext" : { }
  }
},
"status" : {
  "availableReplicas" : 1,
  "conditions" : [ {
    "lastTransitionTime" : "2018-09-06T03:39:40Z",
    "lastUpdateTime" : "2018-09-06T03:39:40Z",
    "message" : "Deployment has minimum availability.",
    "reason" : "MinimumReplicasAvailable",
    "status" : "True",
    "type" : "Available"
  }, {
    "lastTransitionTime" : "2018-09-06T03:39:32Z",
    "lastUpdateTime" : "2018-09-06T03:39:40Z",
    "message" : "ReplicaSet \"deployment-test-865578b586\" has successfully\nprogressed.",
    "reason" : "NewReplicaSetAvailable",
    "status" : "True",
    "type" : "Progressing"
  } ],
  "observedGeneration" : 1,
```

```
"readyReplicas" : 1,  
"replicas" : 1,  
"updatedReplicas" : 1  
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6 Ingress

5.6.1 删除指定 namespace 下的 ingresses

功能介绍

删除Namespace下所有Ingress。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/extensions/v1beta1/namespaces/{namespace}/ingresses

表 5-3070 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3071 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3072 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3073 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3074 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3075 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3076 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3077 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3078 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "items": null,
  "kind": "IngressList",
  "metadata": {
    "resourceVersion": "5160862",
    "selfLink": "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.2 查询指定 namespace 下的 Ingresses

功能介绍

查询Namespace下所有Ingress的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses

表 5-3079 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3080 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3081 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3082 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.extensions.v1beta1.Ingress objects	Items is the list of Ingress.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-3083 io.k8s.api.extensions.v1beta1.Ingress

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3084 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3085 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3086 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3087 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3088 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3089 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3090 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3091 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3092 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3093 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3094 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3095 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3096 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

表 5-3097 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "items": [ {
    "metadata": {
      "annotations": {
        "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
        "kubernetes.io/elb.ip": "192.168.137.182",
        "kubernetes.io/elb.port": "6071"
      }
    }
  }
]
```

```

    },
    "creationTimestamp" : "2018-09-04T02:16:14Z",
    "generation" : 1,
    "labels" : {
      "app" : "redis",
      "isExternal" : "true",
      "zone" : "data"
    },
    "name" : "redis",
    "namespace" : "namespace-test",
    "resourceVersion" : "5161128",
    "selfLink" : "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses/redis",
    "uid" : "7f86c310-afe8-11e8-b6ef-f898ef6c78b4"
  },
  "spec" : {
    "rules" : [ {
      "http" : {
        "paths" : [ {
          "backend" : {
            "serviceName" : "redis",
            "servicePort" : 8080
          },
          "path" : "/"
        } ]
      }
    } ]
  },
  "status" : {
    "loadBalancer" : {
      "ingress" : [ {
        "ip" : "192.168.137.182"
      } ]
    }
  }
},
"kind" : "IngressList",
"metadata" : {
  "resourceVersion" : "5161998",
  "selfLink" : "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses"
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid

状态码	描述
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.3 创建 Ingress

功能介绍

创建Ingress，使用http协议，关联的后端Service为“redis:8080”，使用ELB作为Ingress控制器，ELB的ip为192.168.137.182，端口号为6071。

说明：

若需要在CCI工作负载详情页面的“访问方式”页签中显示对应的Ingress资源，则需要给创建的Ingress资源对象添加labels标签。添加的标签需满足如下要求：

- service的labels中设置的标签必须和负载的selector中matchLabels设置的label一致。

例如，负载中matchLabels标签设置如下：

```
"spec": {
  "replicas": 1,
  "selector": {
    "matchLabels": {
      "app": "redis"
    }
  }
}
```

service中的labels也必须设置为**"app": "redis"**：

```
"metadata": {
  "name": "redis",
  "labels": {
    "app": "redis"
  }
}
```

- ingress中定义的serviceName必须和service中定义的名称一致。

调用方法

请参见[如何调用API](#)。

URI

POST /apis/extensions/v1beta1/namespaces/{namespace}/ingresses

表 5-3098 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3099 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3100 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3101 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3102 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3103 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3104 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3105 io.k8s.api.extensions.v1beta1.IngressSpec

参数	是否必选	参数类型	描述
backend	否	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.

参数	是否必选	参数类型	描述
ingressClassName	否	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	否	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	否	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3106 io.k8s.api.extensions.v1beta1.IngressRule

参数	是否必选	参数类型	描述
host	否	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The <code>:</code> delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly <code>:80</code> for http and <code>:443</code> for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). The wildcard character <code>*</code> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>

参数	是否必选	参数类型	描述
http	否	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.

表 5-3107 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	是否必选	参数类型	描述
paths	是	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	A collection of paths that map requests to backends.

表 5-3108 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	是否必选	参数类型	描述
backend	是	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	否	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.

参数	是否必选	参数类型	描述
pathType	否	String	<p>PathType determines the interpretation of the Path matching. PathType can be one of the following values: *</p> <ul style="list-style-type: none">Exact: Matches the URL path exactly.Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers to the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	否	Map<String,String>	Extension property on the path.

表 5-3109 io.k8s.api.extensions.v1beta1.IngressBackend

参数	是否必选	参数类型	描述
resource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.

参数	是否必选	参数类型	描述
serviceName	否	String	Specifies the name of the referenced service.
servicePort	否	String	Specifies the port of the referenced service.

表 5-3110 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-3111 io.k8s.api.extensions.v1beta1.IngressTLS

参数	是否必选	参数类型	描述
hosts	否	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.

参数	是否必选	参数类型	描述
secretName	否	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3112 io.k8s.api.extensions.v1beta1.IngressStatus

参数	是否必选	参数类型	描述
loadBalancer	否	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3113 io.k8s.api.core.v1.LoadBalancerStatus

参数	是否必选	参数类型	描述
ingress	否	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3114 io.k8s.api.core.v1.LoadBalancerIngress

参数	是否必选	参数类型	描述
hostname	否	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	否	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

响应参数

状态码： 200

表 5-3115 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3116 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3117 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3118 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3119 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3120 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> ' must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3121 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3122 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3123 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3124 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3125 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3126 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3127 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3128 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

状态码: 201

表 5-3129 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3130 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3131 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3132 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3133 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3134 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> ' must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3135 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3136 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3137 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3138 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3139 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3140 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3141 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3142 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

状态码: 202

表 5-3143 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3144 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3145 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3146 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3147 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3148 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> ' must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3149 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3150 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3151 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3152 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3153 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3154 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3155 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3156 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.ip": "192.168.137.182",
      "kubernetes.io/elb.port": "6071"
    },
    "labels": {
      "app": "redis"
    },
    "name": "redis"
  },
  "spec": {
    "rules": [ {
      "http": {
        "paths": [ {
          "backend": {
            "serviceName": "redis",
            "servicePort": 8080
          },
          "path": "/"
        } ]
      }
    } ]
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.port": "6071"
    },
    "creationTimestamp": "2018-09-04T02:16:14Z",
    "generation": 1,
    "labels": {
      "app": "redis",
      "isExternal": "true",
      "zone": "data"
    },
    "name": "redis",
    "namespace": "namespace-test",
    "resourceVersion": "5161127",
    "selfLink": "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses/redis",
    "uid": "7f86c310-afe8-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "rules": [ {
      "http": {
        "paths": [ {
          "backend": {
            "serviceName": "redis",
            "servicePort": 8080
          },
          "path": "/"
        } ]
      }
    } ]
  },
  "status": {
    "loadBalancer": { }
  }
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.4 删除 Ingress

功能介绍

删除Ingress。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}

表 5-3157 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Ingress
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3158 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3159 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3160 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3161 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3162 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.

参数	参数类型	描述
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3163 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.

参数	参数类型	描述
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3164 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3165 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-3166 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3167 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3168 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3169 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "code": 200,
  "details": {
    "group": "extensions",
    "kind": "ingresses",
    "name": "redis",
    "uid": "fa35aa94-afe2-11e8-b6ef-f898ef6c78b4"
  },
  "kind": "Status",
  "metadata": { },
  "status": "Success"
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized

状态码	描述
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.5 查询 Ingress

功能介绍

查询Ingress的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}

表 5-3170 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Ingress
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3171 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3172 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3173 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3174 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3175 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3176 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3177 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3178 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3179 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3180 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	<p>PathType determines the interpretation of the Path matching. PathType can be one of the following values:</p> <ul style="list-style-type: none">* Exact: Matches the URL path exactly.* Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3181 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3182 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3183 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3184 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3185 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3186 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.ip": "192.168.137.182",
      "kubernetes.io/elb.port": "6071"
    },
    "creationTimestamp": "2018-09-04T02:16:14Z",
    "generation": 1,
    "labels": {
      "app": "redis",
      "isExternal": "true",
      "zone": "data"
    },
    "name": "redis",
    "namespace": "namespace-test",
    "resourceVersion": "5161128",
    "selfLink": "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses/redis",
    "uid": "7f86c310-afe8-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "rules": [ {
      "http": {
```



```
"paths" : [ {  
  "backend" : {  
    "serviceName" : "redis",  
    "servicePort" : 8080  
  },  
  "path" : "/"  
}]  
}  
}]  
},  
"status" : {  
  "loadBalancer" : {  
    "ingress" : [ {  
      "ip" : "192.168.137.182"  
    }]  
  }  
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.6 更新 Ingress

功能介绍

更新Ingress。

其中以下字段支持更新：

- metadata.labels
- metadata.generateName
- metadata.annotations
- spec.rules

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}

表 5-3187 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Ingress
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3188 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).

参数	是否必选	参数类型	描述
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3189 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-3190 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-3191 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3192 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3193 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3194 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3195 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3196 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3197 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3198 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3199 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3200 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3201 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3202 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3203 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3204 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

更新Ingress的servicePort值为"8082"。

```
{
  "spec": {
    "rules": [ {
      "http": {
        "paths": [ {
          "backend": {
            "serviceName": "redis",
            "servicePort": 8082
          },
          "path": "/"
        } ]
      }
    } ]
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.ip": "192.168.137.182",
      "kubernetes.io/elb.port": "6071"
    },
    "creationTimestamp": "2018-09-04T02:16:14Z",
  }
}
```



```
"generation" : 3,
"labels" : {
  "app" : "redis",
  "isExternal" : "true",
  "zone" : "data"
},
"name" : "redis",
"namespace" : "namespace-test",
"resourceVersion" : "5165479",
"selfLink" : "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses/redis",
"uid" : "7f86c310-afe8-11e8-b6ef-f898ef6c78b4"
},
"spec" : {
  "rules" : [ {
    "http" : {
      "paths" : [ {
        "backend" : {
          "serviceName" : "redis",
          "servicePort" : 8082
        },
        "path" : "/"
      } ]
    }
  } ]
},
"status" : {
  "loadBalancer" : {
    "ingress" : [ {
      "ip" : "192.168.137.182"
    } ]
  }
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.7 替换 Ingress

功能介绍

替换Ingress。

其中以下字段支持更新：

- metadata.labels
- metadata.generateName
- metadata.annotations
- spec.rules

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}

表 5-3205 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Ingress
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3206 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3207 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3208 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	是否必选	参数类型	描述
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3209 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	是否必选	参数类型	描述
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3210 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3211 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3212 io.k8s.api.extensions.v1beta1.IngressSpec

参数	是否必选	参数类型	描述
backend	否	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.

参数	是否必选	参数类型	描述
ingressClassName	否	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	否	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	否	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3213 io.k8s.api.extensions.v1beta1.IngressRule

参数	是否必选	参数类型	描述
host	否	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The <code>:</code> delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly <code>:80</code> for http and <code>:443</code> for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). The wildcard character <code>*</code> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>

参数	是否必选	参数类型	描述
http	否	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.

表 5-3214 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	是否必选	参数类型	描述
paths	是	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	A collection of paths that map requests to backends.

表 5-3215 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	是否必选	参数类型	描述
backend	是	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	否	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.

参数	是否必选	参数类型	描述
pathType	否	String	<p>PathType determines the interpretation of the Path matching. PathType can be one of the following values: *</p> <ul style="list-style-type: none">Exact: Matches the URL path exactly.Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	否	Map<String,String>	Extension property on the path.

表 5-3216 io.k8s.api.extensions.v1beta1.IngressBackend

参数	是否必选	参数类型	描述
resource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.

参数	是否必选	参数类型	描述
serviceName	否	String	Specifies the name of the referenced service.
servicePort	否	String	Specifies the port of the referenced service.

表 5-3217 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-3218 io.k8s.api.extensions.v1beta1.IngressTLS

参数	是否必选	参数类型	描述
hosts	否	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.

参数	是否必选	参数类型	描述
secretName	否	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3219 io.k8s.api.extensions.v1beta1.IngressStatus

参数	是否必选	参数类型	描述
loadBalancer	否	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3220 io.k8s.api.core.v1.LoadBalancerStatus

参数	是否必选	参数类型	描述
ingress	否	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3221 io.k8s.api.core.v1.LoadBalancerIngress

参数	是否必选	参数类型	描述
hostname	否	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	否	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

响应参数

状态码： 200

表 5-3222 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3223 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3224 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3225 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3226 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3227 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3228 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	A collection of paths that map requests to backends.

表 5-3229 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3230 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3231 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3232 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3233 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3234 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3235 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

状态码: 201

表 5-3236 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3237 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3238 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3239 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3240 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3241 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3242 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3243 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	<p>PathType determines the interpretation of the Path matching. PathType can be one of the following values: * Exact: Matches the URL path exactly. * Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz).</p> <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3244 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3245 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3246 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3247 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3248 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3249 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

将Ingress的servicePort由"8080"替换为"8081"。

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.ip": "192.168.137.182",
      "kubernetes.io/elb.port": "6071"
    },
    "labels": {
      "app": "redis",
      "isExternal": "true",
      "zone": "data"
    },
    "name": "redis"
  },
  "spec": {
    "rules": [ {
      "http": {
        "paths": [ {
          "backend": {
            "serviceName": "redis",
            "servicePort": 8081
          },
          "path": "/"
        } ]
      }
    } ]
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.ip": "192.168.137.182",
      "kubernetes.io/elb.port": "6071"
    },
    "creationTimestamp": "2018-09-04T02:16:14Z",
    "generation": 2,
    "labels": {
      "app": "redis",
      "isExternal": "true",
      "zone": "data"
    },
    "name": "redis",
    "namespace": "namespace-test",
    "resourceVersion": "5162744",
    "selfLink": "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses/redis",
    "uid": "7f86c310-afe8-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "rules": [ {
      "http": {
        "paths": [ {
          "backend": {
            "serviceName": "redis",
            "servicePort": 8081
          },
          "path": "/"
        } ]
      }
    } ]
  },
  "status": {
    "loadBalancer": {
      "ingress": [ {
        "ip": "192.168.137.182"
      } ]
    }
  }
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.6.8 查询 Ingress 状态

功能介绍

查询Ingress状态。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}/status

表 5-3250 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Ingress
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3251 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3252 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3253 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.extensions.v1beta1.IngressSpec object	Spec is the desired state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.extensions.v1beta1.IngressStatus object	Status is the current state of the Ingress. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3254 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3255 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3256 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3257 io.k8s.api.extensions.v1beta1.IngressSpec

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	A default backend capable of servicing requests that don't match any rule. At least one of 'backend' or 'rules' must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
ingressClassName	String	IngressClassName is the name of the IngressClass cluster resource. The associated IngressClass defines which controller will implement the resource. This replaces the deprecated <i>kubernetes.io/ingress.class</i> annotation. For backwards compatibility, when that annotation is set, it must be given precedence over this field. The controller may emit a warning if the field and annotation have different values. Implementations of this API should ignore Ingresses without a class specified. An IngressClass resource may be marked as default, which can be used to set a default value for this field. For more information, refer to the IngressClass documentation.
rules	Array of io.k8s.api.extensions.v1beta1.IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	Array of io.k8s.api.extensions.v1beta1.IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 5-3258 io.k8s.api.extensions.v1beta1.IngressRule

参数	参数类型	描述
host	String	<p>Host is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in RFC 3986: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue.</p> <p>Host can be "precise" which is a domain name without the terminating dot of a network host (e.g. "foo.bar.com") or "wildcard", which is a domain name prefixed with a single wildcard label (e.g. ".foo.com"). <i>The wildcard character</i> must appear by itself as the first DNS label and matches only a single label. You cannot have a wildcard label by itself (e.g. Host == "*"). Requests will be matched against the Host field in the following way: 1. If Host is precise, the request matches this rule if the http host header is equal to Host. 2. If Host is a wildcard, then the request matches this rule if the http host header is to equal to the suffix (removing the first label) of the wildcard rule.</p>
http	io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue object	<p>HTTPIngressRuleValue is a list of http selectors pointing to backends. In the example: http:///? -> backend where where parts of the url correspond to RFC 3986, this resource will be used to match against everything after the last '/' and before the first '?' or '#'.</p>

表 5-3259 io.k8s.api.extensions.v1beta1.HTTPIngressRuleValue

参数	参数类型	描述
paths	Array of io.k8s.api.extensions.v1beta1.HTTPIngressPath objects	<p>A collection of paths that map requests to backends.</p>

表 5-3260 io.k8s.api.extensions.v1beta1.HTTPIngressPath

参数	参数类型	描述
backend	io.k8s.api.extensions.v1beta1.IngressBackend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to.
path	String	Path is matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. When unspecified, all paths from incoming requests are matched.
pathType	String	<p>PathType determines the interpretation of the Path matching. PathType can be one of the following values:</p> <ul style="list-style-type: none">* Exact: Matches the URL path exactly.* Prefix: Matches based on a URL path prefix split by '/'. Matching is done on a path element by element basis. A path element refers is the list of labels in the path split by the '/' separator. A request is a match for path p if every p is an element-wise prefix of p of the request path. Note that if the last element of the path is a substring of the last element in request path, it is not a match (e.g. /foo/bar matches /foo/bar/baz, but does not match /foo/barbaz). <ul style="list-style-type: none">ImplementationSpecific: Interpretation of the Path matching is up to the IngressClass. Implementations can treat this as a separate PathType or treat it identically to Prefix or Exact path types. Implementations are required to support all path types. Defaults to ImplementationSpecific.
property	Map<String,String>	Extension property on the path.

表 5-3261 io.k8s.api.extensions.v1beta1.IngressBackend

参数	参数类型	描述
resource	io.k8s.api.core.v1.TypedLocalObjectReference object	Resource is an ObjectRef to another Kubernetes resource in the namespace of the Ingress object. If resource is specified, serviceName and servicePort must not be specified.
serviceName	String	Specifies the name of the referenced service.

参数	参数类型	描述
servicePort	String	Specifies the port of the referenced service.

表 5-3262 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3263 io.k8s.api.extensions.v1beta1.IngressTLS

参数	参数类型	描述
hosts	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing.

表 5-3264 io.k8s.api.extensions.v1beta1.IngressStatus

参数	参数类型	描述
loadBalancer	io.k8s.api.core.v1.LoadBalancerStatus object	LoadBalancer contains the current status of the load-balancer.

表 5-3265 io.k8s.api.core.v1.LoadBalancerStatus

参数	参数类型	描述
ingress	Array of io.k8s.api.core.v1.LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 5-3266 io.k8s.api.core.v1.LoadBalancerIngress

参数	参数类型	描述
hostname	String	Hostname is set for load-balancer ingress points that are DNS based (typically AWS load-balancers)
ip	String	IP is set for load-balancer ingress points that are IP based (typically GCE or OpenStack load-balancers)

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "extensions/v1beta1",
  "kind": "Ingress",
  "metadata": {
    "annotations": {
      "kubernetes.io/elb.id": "2d48d034-6046-48db-8bb2-53c67e8148b5",
      "kubernetes.io/elb.ip": "192.168.137.182",
      "kubernetes.io/elb.port": "6071"
    },
    "creationTimestamp": "2018-09-04T02:16:14Z",
    "generation": 1,
    "labels": {
      "app": "redis",
      "isExternal": "true",
      "zone": "data"
    },
    "name": "redis",
    "namespace": "namespace-test",
    "resourceVersion": "5161128",
    "selfLink": "/apis/extensions/v1beta1/namespaces/namespace-test/ingresses/redis",
    "uid": "7f86c310-afe8-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "rules": [ {
      "http": {
```

```
    "paths" : [ {  
      "backend" : {  
        "serviceName" : "redis",  
        "servicePort" : 8080  
      },  
      "path" : "/"  
    } ]  
  } ]  
},  
"status" : {  
  "loadBalancer" : {  
    "ingress" : [ {  
      "ip" : "192.168.137.182"  
    } ]  
  }  
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.7 OpenAPIv2

5.7.1 查询 open api swagger 信息

功能介绍

查询open api swagger信息。

调用方法

请参见[如何调用API](#)。

URI

GET /openapi/v2

请求参数

表 5-3267 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

无

请求示例

无

响应示例

无

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.8 VolcanoJob

5.8.1 删除指定 namespace 下的 Volcano Jobs

功能介绍

删除命名空间下的所有Volcano Job。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs

表 5-3268 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3269 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3270 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

响应参数

状态码： 200

表 5-3271 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.

参数	参数类型	描述
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3272 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3273 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3274 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "items": [ {
    "apiVersion": "batch.volcano.sh/v1alpha1",
    "kind": "Job",
    "metadata": {
      "creationTimestamp": "2019-06-26T03:16:26Z",
      "generation": 1,
```

```
"labels" : {
  "app" : "patchlabel"
},
"name" : "openmpi-hello-2-com",
"namespace" : "cci-namespace-42263891",
"resourceVersion" : "7695210",
"selfLink" : "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs/openmpi-
hello-2-com",
"uid" : "c84d86f0-97c0-11e9-9d09-dc9914fb58e0"
},
"spec" : {
  "minAvailable" : 1,
  "plugins" : {
    "env" : [ ],
    "ssh" : [ ],
    "svc" : [ ]
  },
  "queue" : "default",
  "schedulerName" : "volcano",
  "tasks" : [ {
    "name" : "mpimaster",
    "policies" : [ {
      "action" : "CompleteJob",
      "event" : "TaskCompleted"
    } ],
    "replicas" : 1,
    "template" : {
      "spec" : {
        "containers" : [ {
          "command" : [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n'
\\,\\\";\\nmkdir -p /var/run/ssh; /usr/sbin/ssh;\\nmpexec --allow-run-as-root --host ${MPI_HOST} -np 2
mpi_hello_world 003e /home/re" ],
          "image" : "*/*:20202/l00427178/openmpi-hello:3.28",
          "name" : "mpimaster",
          "ports" : [ {
            "containerPort" : 22,
            "name" : "mpijob-port"
          } ],
          "resources" : {
            "limits" : {
              "cpu" : "250m",
              "memory" : "1Gi"
            },
            "requests" : {
              "cpu" : "250m",
              "memory" : "1Gi"
            }
          },
          "workingDir" : "/home"
        } ],
        "imagePullSecrets" : [ {
          "name" : "imagepull-secret"
        } ],
        "restartPolicy" : "OnFailure"
      }
    }
  }, {
    "name" : "mpiworker",
    "replicas" : 2,
    "template" : {
      "spec" : {
        "containers" : [ {
          "command" : [ "/bin/sh", "-c", "mkdir -p /var/run/ssh; /usr/sbin/ssh -D;" ],
          "image" : "*/*:20202/l00427178/openmpi-hello:3.28",
          "name" : "mpiworker",
          "ports" : [ {
            "containerPort" : 22,
            "name" : "mpijob-port"
          } ],
        } ],
      }
    }
  }
], {
```

```
    "resources": {
      "limits": {
        "cpu": "250m",
        "memory": "1Gi"
      },
      "requests": {
        "cpu": "250m",
        "memory": "1Gi"
      }
    },
    "workingDir": "/home"
  }],
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "restartPolicy": "OnFailure"
}
}
}]]
},
"status": {
  "minAvailable": 1,
  "pending": 3,
  "state": {
    "phase": "Inqueue"
  }
}
}],
"kind": "JobList",
"metadata": {
  "continue": "",
  "resourceVersion": "7732232",
  "selfLink": "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable

状态码	描述
504	ServerTimeout

5.8.2 查询指定 namespace 下的 Volcano Jobs

功能介绍

查询命名空间下所有的Volcano Job。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs

表 5-3275 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3276 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3277 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3278 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of sh.volcano.batch.v1alpha1.Job objects	List of jobs. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3279 sh.volcano.batch.v1alpha1.Job

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3280 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3281 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.

参数	参数类型	描述
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3282 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3283 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3284 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.

参数	参数类型	描述
volumeClaim Name	String	The name of the volume claim.

表 5-3285 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3286 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3287 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3288 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3289 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3290 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3291 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3292 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3293 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3294 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3295 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3296 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3297 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

表 5-3298 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "items": [ {
    "apiVersion": "batch.volcano.sh/v1alpha1",
    "kind": "Job",
    "metadata": {
      "creationTimestamp": "2019-06-26T03:16:26Z",
      "generation": 1,
      "name": "openmpi-hello-2-com",
      "namespace": "cci-namespace-42263891",
      "resourceVersion": "7625538",
      "selfLink": "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs/openmpi-hello-2-com",
      "uid": "c84d86f0-97c0-11e9-9d09-dc9914fb58e0"
    },
    "spec": {
      "minAvailable": 1,
      "plugins": {
        "env": [ ],
        "ssh": [ ],
        "svc": [ ]
      },
      "queue": "default",
      "schedulerName": "volcano",
      "tasks": [ {
        "name": "mpimaster",
        "policies": [ {
          "action": "CompleteJob",
          "event": "TaskCompleted"
        } ],
        "replicas": 1,
        "template": {
          "spec": {
            "containers": [ {
              "command": [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n' '\\, \\';\nmkdir -p /var/run/sshdir; /usr/sbin/sshdir;\nmpixexec --allow-run-as-root --host ${MPI_HOST} -np 2 mpi_hello_world 003e /home/re" ],
              "image": "*/.5.235:20202/swr/openmpi-hello:3.28",
              "name": "mpimaster",
              "ports": [ {
                "containerPort": 22,
                "name": "mpijob-port"
              } ],
              "resources": {
                "limits": {
                  "cpu": "250m",
                  "memory": "1Gi"
                },
                "requests": {
                  "cpu": "250m",
                  "memory": "1Gi"
                }
              },
              "workingDir": "/home"
            } ],
            "imagePullSecrets": [ {
              "name": "imagepull-secret"
            } ],
          }
        }
      } ],
    }
  ]
}
```



```

    "restartPolicy" : "OnFailure"
  }
}
}, {
  "name" : "mpiworker",
  "replicas" : 2,
  "template" : {
    "spec" : {
      "containers" : [ {
        "command" : [ "/bin/sh", "-c", "mkdir -p /var/run/ssh; /usr/sbin/ssh -D;" ],
        "image" : "*/*/20202/swr/openmpi-hello:3.28",
        "name" : "mpiworker",
        "ports" : [ {
          "containerPort" : 22,
          "name" : "mpijob-port"
        } ],
        "resources" : {
          "limits" : {
            "cpu" : "250m",
            "memory" : "1Gi"
          },
          "requests" : {
            "cpu" : "250m",
            "memory" : "1Gi"
          }
        },
        "workingDir" : "/home"
      } ],
      "imagePullSecrets" : [ {
        "name" : "imagepull-secret"
      } ],
      "restartPolicy" : "OnFailure"
    }
  }
}
}],
"status" : {
  "minAvailable" : 1,
  "pending" : 3,
  "state" : {
    "phase" : "Inqueue"
  }
}
}],
"kind" : "JobList",
"metadata" : {
  "continue" : "",
  "resourceVersion" : "7678090",
  "selfLink" : "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs"
}
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.8.3 创建 Volcano Job

功能介绍

创建Volcano Job。

调用方法

请参见[如何调用API](#)。

URI

POST /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs

表 5-3299 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3300 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3301 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3302 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	否	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	否	status object	Current status of Job

表 5-3303 spec

参数	是否必选	参数类型	描述
maxRetry	否	Integer	The limit for retrying submitting job, default is 3
minAvailable	否	Integer	The minimal available pods to run for this Job
plugins	否	Object	Enabled task plugins when creating job.

参数	是否必选	参数类型	描述
policies	否	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	否	String	The name of the queue on which job should be created
schedulerName	否	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	否	Array of spec.tasks objects	Tasks specifies the task specification of Job
volumes	否	Array of spec.volumes objects	The volumes for Job

表 5-3304 spec.policies

参数	是否必选	参数类型	描述
action	否	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	否	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	否	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3305 spec.tasks

参数	是否必选	参数类型	描述
name	否	String	Name specifies the name of tasks
policies	否	Array of spec.tasks.policies objects	Specifies the lifecycle of task

参数	是否必选	参数类型	描述
replicas	否	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	否	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3306 spec.tasks.policies

参数	是否必选	参数类型	描述
action	否	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	否	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	否	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3307 spec.volumes

参数	是否必选	参数类型	描述
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain '!'.
volumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	否	String	The name of the volume claim.

表 5-3308 io.k8s.api.core.v1.PersistentVolumeClaim

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	否	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3309 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3310 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3311 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3312 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3313 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-3314 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3315 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3316 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3317 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	否	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	否	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	否	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3318 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	是	String	status is the status of the condition.
type	是	String	type is the type of the condition.

表 5-3319 status

参数	是否必选	参数类型	描述
ControlledResources	否	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	否	Integer	The number of pods which reached phase Succeeded.
failed	否	Integer	The number of pods which reached phase Failed.
minAvailable	否	Integer	The minimal available pods to run for this Job
pending	否	Integer	The number of pending pods.
retryCount	否	Integer	The number that volcano retried to submit the job.
running	否	Integer	The number of running pods.
state	否	status.state object	Current state of Job.
version	否	Integer	Job's current version

表 5-3320 status.state

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
phase	否	String	The phase of Job
reason	否	String	Unique, one-word, CamelCase reason for the condition's last transition.

响应参数

状态码： 200

表 5-3321 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3322 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3323 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3324 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task

参数	参数类型	描述
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3325 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3326 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3327 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3328 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3329 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3330 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3331 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3332 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3333 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3334 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3335 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3336 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3337 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3338 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3339 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

状态码： 201

表 5-3340 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3341 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job

参数	参数类型	描述
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3342 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3343 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3344 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.

参数	参数类型	描述
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3345 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain ':'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3346 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

参数	参数类型	描述
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3347 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3348 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3349 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3350 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3351 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3352 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3353 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3354 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3355 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3356 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3357 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3358 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

状态码： 202

表 5-3359 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3360 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job

参数	参数类型	描述
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3361 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3362 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3363 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.

参数	参数类型	描述
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3364 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain ':'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3365 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

参数	参数类型	描述
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3366 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3367 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3368 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3369 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3370 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3371 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3372 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3373 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3374 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3375 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3376 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3377 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

请求示例

创建Volcano Job。

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "kind": "Job",
  "metadata": {
    "name": "openmpi-hello-2-com"
  },
  "spec": {
    "minAvailable": 1,
    "plugins": {
      "env": [],
      "ssh": [],
      "svc": []
    },
    "schedulerName": "volcano",
    "tasks": [ {
      "name": "mpimaster",
      "policies": [ {
        "action": "CompleteJob",
        "event": "TaskCompleted"
      } ],
      "replicas": 1,
      "template": {
        "spec": {
          "containers": [ {
            "command": [ "/bin/sh", "-c", "MPI_HOST=`cat/etc/volcano/mpiworker.host|tr '\\n' '\\';`nmkdir-
p/var/run/sshhd;/usr/sbin/sshhd;\nmpiexec--allow-run-as-root--host${MPI_HOST}np2mpi_hello_world>/home/re\n" ],
            "image": "20202/swr/openmpi-hello:3.28",
            "name": "mpimaster",
            "ports": [ {
              "containerPort": 22,
              "name": "mpijob-port"
            } ],
            "resources": {
              "limits": {
                "cpu": "250m",
                "memory": "1Gi"
              },
              "requests": {
                "cpu": "250m",
                "memory": "1Gi"
              }
            },
            "workingDir": "/home"
          } ],
          "imagePullSecrets": [ {
            "name": "imagepull-secret"
          } ],
          "restartPolicy": "OnFailure"
        }
      }
    }, {
      "name": "mpiworker",
      "replicas": 2,
      "template": {
        "spec": {
          "containers": [ {
            "command": [ "/bin/sh", "-c", "mkdir-p/var/run/sshhd;/usr/sbin/sshhd-D;" ],
            "image": "20202/swr/openmpi-hello:3.28",
            "name": "mpiworker",
            "ports": [ {
              "containerPort": 22,
              "name": "mpijob-port"
            } ],
            "resources": {
```

```
    "limits": {
      "cpu": "250m",
      "memory": "1Gi"
    },
    "requests": {
      "cpu": "250m",
      "memory": "1Gi"
    }
  },
  "workingDir": "/home"
}],
"imagePullSecrets": [ {
  "name": "imagepull-secret"
}],
"restartPolicy": "OnFailure"
}
}
}]]
}
```

响应示例

状态码: 201

Created

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2019-06-26T06:24:50Z",
    "generation": 1,
    "name": "openmpi-hello-3-com",
    "namespace": "cci-namespace-42263891",
    "resourceVersion": "7681331",
    "selfLink": "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs/openmpi-hello-3-com",
    "uid": "1a32a8c4-97db-11e9-9d09-dc9914fb58e0"
  },
  "spec": {
    "minAvailable": 1,
    "plugins": {
      "env": [ ],
      "ssh": [ ],
      "svc": [ ]
    },
    "queue": "default",
    "schedulerName": "volcano",
    "tasks": [ {
      "name": "mpimaster",
      "policies": [ {
        "action": "CompleteJob",
        "event": "TaskCompleted"
      } ],
      "replicas": 1,
      "template": {
        "spec": {
          "containers": [ {
            "command": [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n' '\\\\n' \\\";\\\";\\nmkdir -p /var/run/ssh; /usr/sbin/sshd; \\nmpiexec --allow-run-as-root --host ${MPI_HOST} -np 2 mpi_hello_world 003e /home/re" ],
            "image": "**:*:20202/swr/openmpi-hello:3.28",
            "name": "mpimaster",
            "ports": [ {
              "containerPort": 22,
              "name": "mpijob-port"
            } ],
            "resources": {
```

```
    "limits" : {
      "cpu" : "250m",
      "memory" : "1Gi"
    },
    "requests" : {
      "cpu" : "250m",
      "memory" : "1Gi"
    }
  },
  "workingDir" : "/home"
}],
"imagePullSecrets" : [ {
  "name" : "imagepull-secret"
} ],
"restartPolicy" : "OnFailure"
}
}
}, {
  "name" : "mpiworker",
  "replicas" : 2,
  "template" : {
    "spec" : {
      "containers" : [ {
        "command" : [ "/bin/sh", "-c", "mkdir -p /var/run/sshd; /usr/sbin/sshd -D;" ],
        "image" : "*/*/20202/swr/openmpi-hello:3.28",
        "name" : "mpiworker",
        "ports" : [ {
          "containerPort" : 22,
          "name" : "mpijob-port"
        } ],
        "resources" : {
          "limits" : {
            "cpu" : "250m",
            "memory" : "1Gi"
          },
          "requests" : {
            "cpu" : "250m",
            "memory" : "1Gi"
          }
        }
      } ],
      "workingDir" : "/home"
    } ],
    "imagePullSecrets" : [ {
      "name" : "imagepull-secret"
    } ],
    "restartPolicy" : "OnFailure"
  }
}
}
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized

状态码	描述
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.8.4 删除 Volcano Job

功能介绍

删除Volcano Job。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs/{name}

表 5-3378 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3379 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3380 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3381 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3382 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3383 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3384 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3385 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3386 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码: 202

表 5-3387 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3388 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3389 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3390 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "details": {
    "group": "batch.volcano.sh",
    "kind": "jobs",
    "name": "openmpi-hello-3-com",
    "uid": "1a32a8c4-97db-11e9-9d09-dc9914fb58e0"
  },
}
```

```
"kind": "Status",  
"metadata": {},  
"status": "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.8.5 查询 Volcano Job 详情

功能介绍

查询Volcano Job的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs/{name}

表 5-3391 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3392 Query 参数

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3393 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3394 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3395 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job

参数	参数类型	描述
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3396 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3397 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3398 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.

参数	参数类型	描述
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3399 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain ':'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3400 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

参数	参数类型	描述
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3401 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3402 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3403 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3404 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3405 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3406 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3407 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3408 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3409 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3410 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3411 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3412 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2019-06-26T06:24:50Z",
    "generation": 1,
    "name": "openmpi-hello-3-com",
    "namespace": "cci-namespace-42263891",
    "resourceVersion": "7681358",
    "selfLink": "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs/openmpi-hello-3-com",
    "uid": "1a32a8c4-97db-11e9-9d09-dc9914fb58e0"
  },
  "spec": {
    "minAvailable": 1,
    "plugins": {
      "env": [],
      "ssh": [],
      "svc": []
    },
    "queue": "default",
    "schedulerName": "volcano",
    "tasks": [ {
      "name": "mpimaster",
      "policies": [ {
        "action": "CompleteJob",
        "event": "TaskCompleted"
      } ],
      "replicas": 1,
      "template": {
        "spec": {
          "containers": [ {
            "command": [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n' '\\,\\\"';\nmkdir -p /var/run/ssh; /usr/sbin/sshd;\nmpixec --allow-run-as-root --host ${MPI_HOST} -np 2\nmpi_hello_world 003e /home/re" ],
            "image": "**:*:20202/swr/openmpi-hello:3.28",
            "name": "mpimaster",
            "ports": [ {
              "containerPort": 22,
              "name": "mpijob-port"
            } ],
            "resources": {
              "limits": {
                "cpu": "250m",
                "memory": "1Gi"
              },
              "requests": {
                "cpu": "250m",
                "memory": "1Gi"
              }
            },
            "workingDir": "/home"
          } ],
          "imagePullSecrets": [ {
            "name": "imagepull-secret"
          } ],
          "restartPolicy": "OnFailure"
        }
      }
    } ]
  }
}
```

```
}
}, {
  "name": "mpiworker",
  "replicas": 2,
  "template": {
    "spec": {
      "containers": [ {
        "command": [ "/bin/sh", "-c", "mkdir -p /var/run/ssh; /usr/sbin/sshd -D;" ],
        "image": "**:*:20202/swr/openmpi-hello:3.28",
        "name": "mpiworker",
        "ports": [ {
          "containerPort": 22,
          "name": "mpijob-port"
        } ],
        "resources": {
          "limits": {
            "cpu": "250m",
            "memory": "1Gi"
          },
          "requests": {
            "cpu": "250m",
            "memory": "1Gi"
          }
        },
        "workingDir": "/home"
      } ],
      "imagePullSecrets": [ {
        "name": "imagepull-secret"
      } ],
      "restartPolicy": "OnFailure"
    }
  }
}
}
}
}, {
  "status": {
    "minAvailable": 1,
    "pending": 3,
    "state": {
      "phase": "Inqueue"
    }
  }
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.8.6 更新 Volcano Job

功能介绍

更新Volcano Job。

The following fields can be updated:

- metadata.labels
- metadata.generateName
- metadata.annotations
- spec.template
- spec.replicas

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs/{name}

表 5-3413 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3414 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3415 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-3416 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-3417 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3418 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job

参数	参数类型	描述
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3419 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3420 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3421 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.

参数	参数类型	描述
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3422 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain ':'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3423 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

参数	参数类型	描述
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3424 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3425 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3426 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3427 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3428 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3429 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3430 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3431 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3432 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3433 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3434 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3435 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

请求示例

更新Volcano Job中的labels值为"app": "patchlabel"。

```
{
  "metadata": {
    "labels": {
      "app": "patchlabel"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2019-06-26T03:16:26Z",
    "generation": 1,
    "labels": {
      "app": "patchlabel"
    },
    "name": "openmpi-hello-2-com",
    "namespace": "cci-namespace-42263891",
    "resourceVersion": "7695210",
    "selfLink": "/apis/batch.volcano.sh/v1alpha1/namespaces/cci-namespace-42263891/jobs/openmpi-hello-2-com",
    "uid": "c84d86f0-97c0-11e9-9d09-dc9914fb58e0"
  },
  "spec": {
    "minAvailable": 1,
    "plugins": {
      "env": [],
      "ssh": [],
      "svc": []
    },
    "queue": "default",
    "schedulerName": "volcano",
    "tasks": [ {
      "name": "mpimaster",
      "policies": [ {
        "action": "CompleteJob",
        "event": "TaskCompleted"
      } ],
      "replicas": 1,
      "template": {
        "spec": {
          "containers": [ {
            "command": [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n\\n' '\\,\\,\\n';\nmkdir -p /var/run/ssh; /usr/sbin/ssh; \nmpixec --allow-run-as-root --host ${MPI_HOST} -np 2\nmpi_hello_world 003e /home/re" ],
            "image": "**:*:*:20202/swr/openmpi-hello:3.28",
            "name": "mpimaster",
            "ports": [ {
              "containerPort": 22,
              "name": "mpijob-port"
            } ],
            "resources": {
              "limits": {
                "cpu": "250m",
                "memory": "1Gi"
              },
              "requests": {
```

```
        "cpu": "250m",
        "memory": "1Gi"
      },
      "workingDir": "/home"
    }],
    "imagePullSecrets": [{
      "name": "imagepull-secret"
    }],
    "restartPolicy": "OnFailure"
  }
}, {
  "name": "mpiworker",
  "replicas": 2,
  "template": {
    "spec": {
      "containers": [{
        "command": ["/bin/sh", "-c", "mkdir -p /var/run/ssh; /usr/sbin/ssh -D;"],
        "image": "*/*:20202/swr/openmpi-hello:3.28",
        "name": "mpiworker",
        "ports": [{
          "containerPort": 22,
          "name": "mpijob-port"
        }],
        "resources": {
          "limits": {
            "cpu": "250m",
            "memory": "1Gi"
          },
          "requests": {
            "cpu": "250m",
            "memory": "1Gi"
          }
        }
      }],
      "workingDir": "/home"
    }],
    "imagePullSecrets": [{
      "name": "imagepull-secret"
    }],
    "restartPolicy": "OnFailure"
  }
}
}],
"status": {
  "minAvailable": 1,
  "pending": 3,
  "state": {
    "phase": "Inqueue"
  }
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.8.7 替换 Volcano Job

功能介绍

替换Volcano Job。

The following fields can be updated:

- metadata.labels
- metadata.generateName
- metadata.annotations
- spec.template
- spec.replicas

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/batch.volcano.sh/v1alpha1/namespaces/{namespace}/jobs/{name}

表 5-3436 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3437 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3438 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3439 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	否	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	否	status object	Current status of Job

表 5-3440 spec

参数	是否必选	参数类型	描述
maxRetry	否	Integer	The limit for retrying submitting job, default is 3
minAvailable	否	Integer	The minimal available pods to run for this Job
plugins	否	Object	Enabled task plugins when creating job.

参数	是否必选	参数类型	描述
policies	否	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	否	String	The name of the queue on which job should be created
schedulerName	否	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	否	Array of spec.tasks objects	Tasks specifies the task specification of Job
volumes	否	Array of spec.volumes objects	The volumes for Job

表 5-3441 spec.policies

参数	是否必选	参数类型	描述
action	否	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	否	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	否	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3442 spec.tasks

参数	是否必选	参数类型	描述
name	否	String	Name specifies the name of tasks
policies	否	Array of spec.tasks.policies objects	Specifies the lifecycle of task

参数	是否必选	参数类型	描述
replicas	否	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	否	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3443 spec.tasks.policies

参数	是否必选	参数类型	描述
action	否	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	否	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	否	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3444 spec.volumes

参数	是否必选	参数类型	描述
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '!'.
volumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	否	String	The name of the volume claim.

表 5-3445 io.k8s.api.core.v1.PersistentVolumeClaim

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	否	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3446 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3447 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3448 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3449 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3450 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-3451 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3452 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3453 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3454 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	否	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	否	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	否	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3455 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	是	String	status is the status of the condition.
type	是	String	type is the type of the condition.

表 5-3456 status

参数	是否必选	参数类型	描述
ControlledResources	否	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	否	Integer	The number of pods which reached phase Succeeded.
failed	否	Integer	The number of pods which reached phase Failed.
minAvailable	否	Integer	The minimal available pods to run for this Job
pending	否	Integer	The number of pending pods.
retryCount	否	Integer	The number that volcano retried to submit the job.
running	否	Integer	The number of running pods.
state	否	status.state object	Current state of Job.
version	否	Integer	Job's current version

表 5-3457 status.state

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
phase	否	String	The phase of Job
reason	否	String	Unique, one-word, CamelCase reason for the condition's last transition.

响应参数

状态码： 200

表 5-3458 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3459 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3460 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3461 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task

参数	参数类型	描述
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3462 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3463 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3464 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3465 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3466 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3467 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3468 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3469 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3470 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3471 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3472 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3473 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3474 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3475 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3476 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

状态码： 201

表 5-3477 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata
spec	spec object	Specification of the desired behavior of a cron job, including the minAvailable
status	status object	Current status of Job

表 5-3478 spec

参数	参数类型	描述
maxRetry	Integer	The limit for retrying submitting job, default is 3
minAvailable	Integer	The minimal available pods to run for this Job
plugins	Object	Enabled task plugins when creating job.
policies	Array of spec.policies objects	Specifies the default lifecycle of tasks
queue	String	The name of the queue on which job should be created
schedulerName	String	SchedulerName is the default value of <i>tasks.template.spec.schedulerName</i> .
tasks	Array of spec.tasks objects	Tasks specifies the task specification of Job

参数	参数类型	描述
volumes	Array of spec.volumes objects	The volumes for Job

表 5-3479 spec.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3480 spec.tasks

参数	参数类型	描述
name	String	Name specifies the name of tasks
policies	Array of spec.tasks.policies objects	Specifies the lifecycle of task
replicas	Integer	Replicas specifies the replicas of this TaskSpec in Job
template	Object	Specifies the pod that will be created for this TaskSpec when executing a Job

表 5-3481 spec.tasks.policies

参数	参数类型	描述
action	String	The action that will be taken to the PodGroup according to Event. One of "Restart", "None". Default to None.
event	String	The Event recorded by scheduler; the controller takes actions according to this Event.

参数	参数类型	描述
timeout	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 5-3482 spec.volumes

参数	参数类型	描述
mountPath	String	Path within the container at which the volume should be mounted. Must not contain ':'.
volumeClaim	io.k8s.api.core.v1.PersistentVolumeClaim object	VolumeClaim defines the PVC used by the VolumeMount.
volumeClaim Name	String	The name of the volume claim.

表 5-3483 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

参数	参数类型	描述
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-3484 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3485 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3486 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3487 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3488 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3489 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3490 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3491 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3492 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-3493 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-3494 status

参数	参数类型	描述
ControlledResources	Map<String,String>	All of the resources that are controlled by this job.
Succeeded	Integer	The number of pods which reached phase Succeeded.
failed	Integer	The number of pods which reached phase Failed.
minAvailable	Integer	The minimal available pods to run for this Job
pending	Integer	The number of pending pods.
retryCount	Integer	The number that volcano retried to submit the job.
running	Integer	The number of running pods.
state	status.state object	Current state of Job.
version	Integer	Job's current version

表 5-3495 status.state

参数	参数类型	描述
message	String	Human-readable message indicating details about last transition.
phase	String	The phase of Job
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.

请求示例

将创建的Volcano Job的其中一个task的replicas从1修改为2。

```
{
  "apiVersion": "batch.volcano.sh/v1alpha1",
  "kind": "Job",
  "metadata": {
    "name": "openmpi-hello-2-com",
    "namespace": "cci-gpu",
    "resourceVersion": "125597961"
  },
  "spec": {
    "minAvailable": 2,
    "plugins": {
      "env": [],
      "ssh": [],
      "svc": []
    },
    "queue": "default",
    "schedulerName": "volcano",
    "tasks": [ {
      "name": "mpimaster",
      "policies": [ {
        "action": "CompleteJob",
        "event": "TaskCompleted"
      } ],
      "replicas": 1,
      "template": {
        "spec": {
          "containers": [ {
            "command": [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n\\n' '\\ \\ \\ \\n\\n\\n\\n';\nmkdir -p /var/run/ssh; /usr/sbin/ssh; \nmkpiexec --allow-run-as-root --host ${MPI_HOST} -np 2
mpi_hello_world 003e /home/re" ],
            "image": "*:*:20202/swr/openmpi-hello:3.28",
            "name": "mpimaster",
            "ports": [ {
              "containerPort": 22,
              "name": "mpijob-port"
            } ],
            "resources": {
              "limits": {
                "cpu": "250m",
                "memory": "1Gi"
              },
              "requests": {
                "cpu": "250m",
                "memory": "1Gi"
              }
            }
          } ],
          "workingDir": "/home"
        } ],
        "imagePullSecrets": [ {
          "name": "imagepull-secret"
        } ],
        "restartPolicy": "OnFailure"
      }
    } ],
    "name": "mpiworker",
    "replicas": 2,
    "template": {
      "spec": {
        "containers": [ {
          "command": [ "/bin/sh", "-c", "mkdir -p /var/run/ssh; /usr/sbin/ssh -D;" ],
          "image": "*:*:20202/swr/openmpi-hello:3.28",
          "name": "mpiworker",
          "ports": [ {
            "containerPort": 22,
```

```
    "name" : "mpijob-port"
  } ],
  "resources" : {
    "limits" : {
      "cpu" : "250m",
      "memory" : "1Gi"
    },
    "requests" : {
      "cpu" : "250m",
      "memory" : "1Gi"
    }
  },
  "workingDir" : "/home"
} ],
"imagePullSecrets" : [ {
  "name" : "imagepull-secret"
} ],
"restartPolicy" : "OnFailure"
}
}
} ]
},
"status" : {
  "minAvailable" : 1,
  "pending" : 3,
  "state" : {
    "phase" : "Succeed"
  }
}
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion" : "batch.volcano.sh/v1alpha1",
  "kind" : "Job",
  "metadata" : {
    "creationTimestamp" : "2019-07-11T08:09:41Z",
    "generation" : 2,
    "name" : "openmpi-hello-2-com",
    "namespace" : "zjh-gpu",
    "resourceVersion" : "125620754",
    "selfLink" : "/apis/batch.volcano.sh/v1alpha1/namespaces/zjh-gpu/jobs/openmpi-hello-2-com",
    "uid" : "3bdba739-a3b3-11e9-8865-501d93440997"
  },
  "spec" : {
    "minAvailable" : 2,
    "plugins" : {
      "env" : [ ],
      "ssh" : [ ],
      "svc" : [ ]
    },
    "queue" : "default",
    "schedulerName" : "volcano",
    "tasks" : [ {
      "name" : "mpimaster",
      "policies" : [ {
        "action" : "CompleteJob",
        "event" : "TaskCompleted"
      } ],
      "replicas" : 1,
      "template" : {
        "spec" : {
          "containers" : [ {
            "command" : [ "/bin/sh", "-c", "MPI_HOST=`cat /etc/volcano/mpiworker.host | tr '\\n'"
```

```
\,\";\nmkdir -p /var/run/ssh; /usr/sbin/ssh; \nmpirun --allow-run-as-root --host ${MPI_HOST} -np 2  
mpi_hello_world 003e /home/re" ],  
  "image": "*/*:20202/swr/openmpi-hello:3.28",  
  "name": "mpimaster",  
  "ports": [ {  
    "containerPort": 22,  
    "name": "mpijob-port"  
  } ],  
  "resources": {  
    "limits": {  
      "cpu": "250m",  
      "memory": "1Gi"  
    },  
    "requests": {  
      "cpu": "250m",  
      "memory": "1Gi"  
    }  
  },  
  "workingDir": "/home"  
},  
  "imagePullSecrets": [ {  
    "name": "imagepull-secret"  
  } ],  
  "restartPolicy": "OnFailure"  
}  
}, {  
  "name": "mpiworker",  
  "replicas": 2,  
  "template": {  
    "spec": {  
      "containers": [ {  
        "command": [ "/bin/sh", "-c", "mkdir -p /var/run/ssh; /usr/sbin/ssh -D;" ],  
        "image": "*/*:20202/swr/openmpi-hello:3.28",  
        "name": "mpiworker",  
        "ports": [ {  
          "containerPort": 22,  
          "name": "mpijob-port"  
        } ],  
        "resources": {  
          "limits": {  
            "cpu": "250m",  
            "memory": "1Gi"  
          },  
          "requests": {  
            "cpu": "250m",  
            "memory": "1Gi"  
          }  
        },  
        "workingDir": "/home"  
      } ],  
      "imagePullSecrets": [ {  
        "name": "imagepull-secret"  
      } ],  
      "restartPolicy": "OnFailure"  
    }  
  }  
},  
  "status": {  
    "minAvailable": 1,  
    "pending": 3,  
    "state": {  
      "phase": "Inqueue"  
    }  
  }  
}  
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.9 Namespace

5.9.1 查询所有 Namespaces

功能介绍

该API用于获取集群中该用户当前项目下所有Namespace的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces

表 5-3496 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3497 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3498 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Namespace objects	Items is the list of Namespace objects in the list. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3499 io.k8s.api.core.v1.Namespace

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.NamespaceSpec object	Spec defines the behavior of the Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.NamespaceStatus object	Status describes the current status of a Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3500 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3501 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3502 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3503 io.k8s.api.core.v1.NamespaceSpec

参数	参数类型	描述
finalizers	Array of strings	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3504 io.k8s.api.core.v1.NamespaceStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.NamespaceCondition objects	Represents the latest available observations of a namespace's current state.
phase	String	Phase is the current lifecycle phase of the namespace. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3505 io.k8s.api.core.v1.NamespaceCondition

参数	参数类型	描述
lastTransitionTime	String	Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.
message	String	
reason	String	
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of namespace controller condition.

表 5-3506 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2017-06-24T02:05:16Z",
      "name": "default",
      "resourceVersion": "6",
      "selfLink": "/api/v1/namespaces/default",
```

```

    "uid" : "90dd5244-5881-11e7-b5d7-fa163e08a2fd"
  },
  "spec" : {
    "finalizers" : [ "kubernetes" ]
  },
  "status" : {
    "phase" : "Active"
  }
}, {
  "metadata" : {
    "creationTimestamp" : "2017-06-24T02:05:17Z",
    "name" : "kube-system",
    "resourceVersion" : "25",
    "selfLink" : "/api/v1/namespaces/kube-system",
    "uid" : "9178fce6-5881-11e7-b5d7-fa163e08a2fd"
  },
  "spec" : {
    "finalizers" : [ "kubernetes" ]
  },
  "status" : {
    "phase" : "Active"
  }
}],
"kind" : "NamespaceList",
"metadata" : {
  "resourceVersion" : "594181",
  "selfLink" : "/api/v1/namespaces"
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.9.2 创建 Namespace

功能介绍

创建一个Namespace。

当前云容器实例提供“通用计算型”和“GPU加速型”两种类型的资源，创建命名空间时需要选择资源类型，后续创建的负载中容器就运行在此类型的集群上。调用接口时必须指定metadata.annotations中namespace.kubernetes.io/flavor字段为如下值。

- general-computing: 通用计算型，基本水平的计算、存储和网络资源，适用于通用工作负载场景。
- gpu-accelerated: GPU加速型，突出的图形计算能力，适用于AI等高性能场景。

Namespace下必须要创建一个Network，用于定义kubernetes中一个namespace内的网络与华为云虚拟私有云服务的子网和vpc的映射关系，具体请参见[Namespace和Network](#)。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces

表 5-3507 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3508 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3509 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	是否必选	参数类型	描述
spec	否	io.k8s.api.core.v1.NamespaceSpec object	Spec defines the behavior of the Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.core.v1.NamespaceStatus object	Status describes the current status of a Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3510 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3511 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3512 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3513 io.k8s.api.core.v1.NamespaceSpec

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3514 io.k8s.api.core.v1.NamespaceStatus

参数	是否必选	参数类型	描述
conditions	否	Array of io.k8s.api.core.v1.NamespaceCondition objects	Represents the latest available observations of a namespace's current state.
phase	否	String	Phase is the current lifecycle phase of the namespace. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3515 io.k8s.api.core.v1.NamespaceCondition

参数	是否必选	参数类型	描述
lastTransitionTime	否	String	Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.
message	否	String	
reason	否	String	
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of namespace controller condition.

响应参数

状态码： 200

表 5-3516 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.NamespaceSpec object	Spec defines the behavior of the Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.NamespaceStatus object	Status describes the current status of a Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3517 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3518 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3519 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3520 io.k8s.api.core.v1.NamespaceSpec

参数	参数类型	描述
finalizers	Array of strings	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3521 io.k8s.api.core.v1.NamespaceStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.NamespaceCondition objects	Represents the latest available observations of a namespace's current state.
phase	String	Phase is the current lifecycle phase of the namespace. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3522 io.k8s.api.core.v1.NamespaceCondition

参数	参数类型	描述
lastTransitionTime	String	Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.
message	String	
reason	String	
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of namespace controller condition.

状态码: 201

表 5-3523 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.NamespaceSpec object	Spec defines the behavior of the Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.NamespaceStatus object	Status describes the current status of a Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3524 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3525 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3526 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3527 io.k8s.api.core.v1.NamespaceSpec

参数	参数类型	描述
finalizers	Array of strings	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3528 io.k8s.api.core.v1.NamespaceStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.NamespaceCondition objects	Represents the latest available observations of a namespace's current state.
phase	String	Phase is the current lifecycle phase of the namespace. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3529 io.k8s.api.core.v1.NamespaceCondition

参数	参数类型	描述
lastTransitionTime	String	Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.
message	String	
reason	String	
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of namespace controller condition.

状态码: 202

表 5-3530 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.NamespaceSpec object	Spec defines the behavior of the Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.NamespaceStatus object	Status describes the current status of a Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3531 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3532 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3533 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3534 io.k8s.api.core.v1.NamespaceSpec

参数	参数类型	描述
finalizers	Array of strings	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3535 io.k8s.api.core.v1.NamespaceStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.NamespaceCondition objects	Represents the latest available observations of a namespace's current state.
phase	String	Phase is the current lifecycle phase of the namespace. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3536 io.k8s.api.core.v1.NamespaceCondition

参数	参数类型	描述
lastTransitionTime	String	Time is a wrapper around time.Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.
message	String	
reason	String	
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of namespace controller condition.

请求示例

创建GPU加速型命名空间。

```
{
  "apiVersion": "v1",
  "kind": "Namespace",
  "metadata": {
    "annotations": {
      "namespace.kubernetes.io/flavor": "gpu-accelerated"
    },
    "labels": {
      "sys_enterprise_project_id": "0"
    },
    "name": "namespace-test"
  }
}
```

响应示例

状态码: 201

Created

```
{
  "metadata": {
    "annotations": {
      "namespace.kubernetes.io/flavor": "gpu-accelerated",
      "pv.kubernetes.io/enable-dynamic-provisioning": "true",
      "tenant.kubernetes.io/domain-id": "aadb43c0b14c4cafbcfff483d075987",
      "tenant.kubernetes.io/domain-name": "cci",
      "tenant.kubernetes.io/project-id": "51bf52609f2a49c68bfda3398817b376",
      "tenant.kubernetes.io/project-name": "southchina"
    },
    "creationTimestamp": "2018-09-03T11:20:48Z",
    "labels": {
      "sys_enterprise_project_id": "0503dda897000fed2f78c00909158a4d"
    },
    "name": "namespace-test",
    "resourceVersion": "5016746",
    "selfLink": "/api/v1/namespaces/namespace-test",
    "uid": "68a68c5a-af6b-11e8-8f17-c81fbc371a17"
  },
  "spec": {
    "finalizers": [ "kubernetes" ]
  },
  "status": {
    "phase": "Active"
  }
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.9.3 删除 Namespace

功能介绍

删除一个Namespace。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{name}

表 5-3537 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Namespace

表 5-3538 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3539 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3540 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3541 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3542 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3543 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3544 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3545 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-3546 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3547 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3548 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3549 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

删除命名空间，删除前持续10S。

```
{  
  "apiVersion": "v1",  
  "gracePeriodSeconds": 10,  
  "kind": "DeleteOptions"  
}
```

响应示例

无

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.9.4 查询 Namespace

功能介绍

查询Namespace的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{name}

表 5-3550 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Namespace

表 5-3551 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3552 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3553 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.NamespaceSpec object	Spec defines the behavior of the Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.NamespaceStatus object	Status describes the current status of a Namespace. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3554 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3555 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3556 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3557 io.k8s.api.core.v1.NamespaceSpec

参数	参数类型	描述
finalizers	Array of strings	Finalizers is an opaque list of values that must be empty to permanently remove object from storage. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3558 io.k8s.api.core.v1.NamespaceStatus

参数	参数类型	描述
conditions	Array of io.k8s.api.core.v1.NamespaceCondition objects	Represents the latest available observations of a namespace's current state.
phase	String	Phase is the current lifecycle phase of the namespace. More info: https://kubernetes.io/docs/tasks/administer-cluster/namespaces/

表 5-3559 io.k8s.api.core.v1.NamespaceCondition

参数	参数类型	描述
lastTransitionTime	String	Time is a wrapper around time. Time which supports correct marshaling to YAML and JSON. Wrappers are provided for many of the factory methods that the time package offers.
message	String	
reason	String	
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of namespace controller condition.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "metadata": {
    "annotations": {
      "namespace.kubernetes.io/flavor": "gpu-accelerated",
      "pv.kubernetes.io/enable-dynamic-provisioning": "true",
      "tenant.kubernetes.io/domain-id": "aadb43c0b14c4cafbccfff483d075987",
      "tenant.kubernetes.io/domain-name": "cci",
      "tenant.kubernetes.io/project-id": "51bf52609f2a49c68bfda3398817b376",
      "tenant.kubernetes.io/project-name": "southchina"
    },
    "creationTimestamp": "2018-09-03T11:20:48Z",
    "name": "namespace-test",
    "resourceVersion": "5016746",
    "selfLink": "/api/v1/namespaces/namespace-test",
  }
}
```

```
"uid" : "68a68c5a-af6b-11e8-8f17-c81fbe371a17"  
},  
"spec" : {  
  "finalizers" : [ "kubernetes" ]  
},  
"status" : {  
  "phase" : "Active"  
}  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.10 ClusterRole

5.10.1 获取 ClusterRole 列表

功能介绍

This API is used to list or watch objects of kind ClusterRole

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io/v1/clusterroles

表 5-3560 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3561 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3562 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.rbac.v1.ClusterRole objects	Items is a list of ClusterRoles
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard object's metadata.

表 5-3563 io.k8s.api.rbac.v1.ClusterRole

参数	参数类型	描述
aggregationRule	io.k8s.api.rbac.v1.AggregationRule object	AggregationRule is an optional field that describes how to build the Rules for this ClusterRole. If AggregationRule is set, then the Rules are controller managed and direct changes to Rules will be stomped by the controller.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
rules	Array of io.k8s.api.rbac.v1.PolicyRule objects	Rules holds all the PolicyRules for this ClusterRole

表 5-3564 io.k8s.api.rbac.v1.AggregationRule

参数	参数类型	描述
clusterRoleSelectors	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector objects	ClusterRoleSelectors holds a list of selectors which will be used to find ClusterRoles and create the rules. If any of the selectors match, then the ClusterRole's permissions will be added

表 5-3565 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3566 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3567 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3568 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3569 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3570 io.k8s.api.rbac.v1.PolicyRule

参数	参数类型	描述
apiGroups	Array of strings	APIGroups is the name of the APIGroup that contains the resources. If multiple API groups are specified, any action requested against one of the enumerated resources in any API group will be allowed.

参数	参数类型	描述
nonResourceURLs	Array of strings	NonResourceURLs is a set of partial urls that a user should have access to. *s are allowed, but only as the full, final step in the path Since non-resource URLs are not namespaced, this field is only applicable for ClusterRoles referenced from a ClusterRoleBinding. Rules can either apply to API resources (such as "pods" or "secrets") or non-resource URL paths (such as "/api"), but not both.
resourceNames	Array of strings	ResourceNames is an optional white list of names that the rule applies to. An empty set means that everything is allowed.
resources	Array of strings	Resources is a list of resources this rule applies to. ResourceAll represents all resources.
verbs	Array of strings	Verbs is a list of Verbs that apply to ALL the ResourceKinds and AttributeRestrictions contained in this rule. VerbAll represents all kinds.

表 5-3571 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2018-11-27T01:55:44Z",
      "name": "secret-reader",
      "resourceVersion": "4619",
      "selfLink": "/apis/rbac.authorization.k8s.io/v1/clusterroles/secret-reader",
```

```
"uid" : "8d358854-f1e7-11e8-b449-fa163ec24e06"
},
"rules" : [ {
  "apiGroups" : [ "" ],
  "resources" : [ "secrets" ],
  "verbs" : [ "get", "watch", "list" ]
} ],
"kind" : "ClusterRoleList",
"metadata" : {
  "resourceVersion" : "4622",
  "selfLink" : "/apis/rbac.authorization.k8s.io/v1/clusterroles"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.10.2 查询指定的 ClusterRole

功能介绍

This API is used to read the specified ClusterRole

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}

表 5-3572 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ClusterRole

表 5-3573 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3574 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3575 响应 Body 参数

参数	参数类型	描述
aggregationRule	io.k8s.api.rbac.v1.AggregationRule object	AggregationRule is an optional field that describes how to build the Rules for this ClusterRole. If AggregationRule is set, then the Rules are controller managed and direct changes to Rules will be stomped by the controller.

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
rules	Array of io.k8s.api.rbac.v1.PolicyRule objects	Rules holds all the PolicyRules for this ClusterRole

表 5-3576 io.k8s.api.rbac.v1.AggregationRule

参数	参数类型	描述
clusterRoleSelectors	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector objects	ClusterRoleSelectors holds a list of selectors which will be used to find ClusterRoles and create the rules. If any of the selectors match, then the ClusterRole's permissions will be added

表 5-3577 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3578 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3579 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3580 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3581 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3582 io.k8s.api.rbac.v1.PolicyRule

参数	参数类型	描述
apiGroups	Array of strings	APIGroups is the name of the APIGroup that contains the resources. If multiple API groups are specified, any action requested against one of the enumerated resources in any API group will be allowed.

参数	参数类型	描述
nonResourceURLs	Array of strings	NonResourceURLs is a set of partial urls that a user should have access to. *s are allowed, but only as the full, final step in the path Since non-resource URLs are not namespaced, this field is only applicable for ClusterRoles referenced from a ClusterRoleBinding. Rules can either apply to API resources (such as "pods" or "secrets") or non-resource URL paths (such as "/api"), but not both.
resourceNames	Array of strings	ResourceNames is an optional white list of names that the rule applies to. An empty set means that everything is allowed.
resources	Array of strings	Resources is a list of resources this rule applies to. ResourceAll represents all resources.
verbs	Array of strings	Verbs is a list of Verbs that apply to ALL the ResourceKinds and AttributeRestrictions contained in this rule. VerbAll represents all kinds.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "kind": "ClusterRole",
  "metadata": {
    "creationTimestamp": "2018-11-27T03:03:00Z",
    "name": "secret-reader",
    "resourceVersion": "13211",
    "selfLink": "/apis/rbac.authorization.k8s.io/v1/clusterroles/secret-reader",
    "uid": "f2cf199e-f1f0-11e8-b449-fa163ec24e06"
  },
  "rules": [ {
    "apiGroups": [ "" ],
    "resources": [ "secrets" ],
    "verbs": [ "get", "watch", "list" ]
  } ]
}
```

状态码

状态码	描述
200	OK

状态码	描述
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.11 Secret

5.11.1 删除指定 namespace 下的 Secrets

功能介绍

删除Namespace下所有Secret。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/secrets

表 5-3583 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3584 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3585 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3586 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3587 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3588 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3589 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3590 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3591 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "data": {
      "data_key_1": "dGVzdA==",
      "data_key_2": "dGVzdA==",
      "data_key_3": "dGVzdA==",
      "data_key_4": "dGVzdA==",
      "data_key_5": "dGVzdA==",
      "data_key_6": "dGVzdA==",
      "data_key_7": "dGVzdA==",
      "data_key_8": "dGVzdA==",
      "data_key_9": "dGVzdA==",
      "data_key_10": "dGVzdA=="
    },
    "metadata": {
      "annotations": {
        "description": "",
        "secret.cci.io/namespace-uid": "c1bbdefd-c577-11e9-9355-dc9914fb58e0"
      },
      "creationTimestamp": "2019-09-03T07:29:57Z",
      "labels": {
        "label_key": "label_value"
      },
      "name": "secret-1",
      "namespace": "testvipchange2",
      "resourceVersion": "39601847",
      "selfLink": "/api/v1/namespaces/testvipchange2/secrets/secret-1",
      "uid": "a1bfe325-ce1c-11e9-83db-dc9914fb58e0"
    }
  }
],
}
```

```
"type" : "Opaque"
}, {
  "data" : {
    "data_key_1" : "dGVzdA==",
    "data_key_2" : "dGVzdA==",
    "data_key_3" : "dGVzdA==",
    "data_key_4" : "dGVzdA==",
    "data_key_5" : "dGVzdA==",
    "data_key_6" : "dGVzdA==",
    "data_key_7" : "dGVzdA==",
    "data_key_8" : "dGVzdA==",
    "data_key_9" : "dGVzdA==",
    "data_key_10" : "dGVzdA=="
  },
  "metadata" : {
    "annotations" : {
      "description" : "",
      "secret.cci.io/namespace-uid" : "c1bbdefd-c577-11e9-9355-dc9914fb58e0"
    },
    "creationTimestamp" : "2019-09-03T07:29:58Z",
    "labels" : {
      "label_key" : "label_value"
    },
    "name" : "secret-3",
    "namespace" : "testvipchange2",
    "resourceVersion" : "39601852",
    "selfLink" : "/api/v1/namespaces/testvipchange2/secrets/secret-3",
    "uid" : "a1cbf6a2-ce1c-11e9-83db-dc9914fb58e0"
  },
  "type" : "Opaque"
}],
"kind" : "SecretList",
"metadata" : {
  "resourceVersion" : "39602472",
  "selfLink" : "/api/v1/namespaces/testvipchange2/secrets"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError

状态码	描述
503	ServiceUnavailable
504	ServerTimeout

5.11.2 查询指定 namespace 下的 Secrets

功能介绍

查询指定namespace下的所有Secret对象。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/secrets

表 5-3592 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3593 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3594 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3595 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Secret objects	Items is a list of secret objects. More info: https://kubernetes.io/docs/concepts/configuration/secret
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3596 io.k8s.api.core.v1.Secret

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '.'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4

参数	参数类型	描述
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3597 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3598 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3599 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3600 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "apiVersion": "v1",
    "data": {
      "KEY": "VkFMVUU="
    },
    "kind": "Secret",
    "metadata": {
      "annotations": {
        "secret.cci.io/namespace-uid": "a9862108-5ef1-4ca5-bd26-8ed19da4ab0e"
      },
      "creationTimestamp": "2022-09-06T01:51:20Z",
      "name": "test-secret",
      "namespace": "namespace-test",
      "resourceVersion": "41271056",
      "selfLink": "/api/v1/namespaces/namespace-test/secrets/test-secret",
      "uid": "81ddca10-6784-40d7-a58e-05f1a6088d94"
    },
    "type": "Opaque"
  } ],
  "kind": "SecretList",
  "metadata": {
    "resourceVersion": "41327545",
    "selfLink": "/api/v1/namespaces/namespace-test/secrets"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests

状态码	描述
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.11.3 创建 Secret

功能介绍

创建Secret, Kubernetes提供了Secret来处理敏感信息。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/secrets

表 5-3601 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3602 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3603 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3604 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	是否必选	参数类型	描述
data	否	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	否	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by <code>ImmutableEphemeralVolumes</code> feature gate.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	否	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.

参数	是否必选	参数类型	描述
type	否	String	Used to facilitate programmatic handling of secret data.

表 5-3605 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	是否必选	参数类型	描述
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3606 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3607 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

响应参数

状态码： 200

表 5-3608 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '.'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3609 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3610 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3611 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

状态码： 201

表 5-3612 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3613 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3614 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3615 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

状态码： 202

表 5-3616 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3617 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3618 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3619 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

创建的Secret中包含两条Key-Value。

```
{
  "apiVersion": "v1",
  "data": {
    "key1": "MWYyZDFIMmU2N2Rm",
    "key2": "YWRtaW4="
  },
  "kind": "Secret",
  "metadata": {
```



```
"name" : "secret-test"  
},  
"type" : "Opaque"  
}
```

响应示例

状态码: 201

Created

```
{  
  "apiVersion" : "v1",  
  "data" : {  
    "key1" : "MWYyZDFlMmU2N2Rm",  
    "key2" : "YWRtaW4="  
  },  
  "kind" : "Secret",  
  "metadata" : {  
    "creationTimestamp" : "2018-09-04T03:59:19Z",  
    "name" : "secret-test",  
    "namespace" : "namespace-test",  
    "resourceVersion" : "5177770",  
    "selfLink" : "/api/v1/namespaces/namespace-test/secrets/secret-test",  
    "uid" : "e6170b6d-aff6-11e8-8f17-c81fbe371a17"  
  },  
  "type" : "Opaque"  
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable

状态码	描述
504	ServerTimeout

5.11.4 删除 Secret

功能介绍

删除Secret。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/secrets/{name}

表 5-3620 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Secret
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3621 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3622 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3623 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3624 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3625 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.

参数	参数类型	描述
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3626 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.

参数	参数类型	描述
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3627 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3628 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-3629 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3630 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3631 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3632 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 10,
  "kind": "DeleteOptions"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "code": 200,
  "details": {
    "kind": "secrets",
    "name": "secret-test",
    "uid": "e6170b6d-aff6-11e8-8f17-c81fbe371a17"
  },
  "kind": "Status",
  "metadata": { },
  "status": "Success"
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest

状态码	描述
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.11.5 查询 Secret

功能介绍

查询Secret的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/secrets/{name}

表 5-3633 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Secret
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3634 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3635 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3636 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '.'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3637 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3638 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3639 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

无

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "v1",
```

```
"data" : {
  "key1" : "MWYyZDFIMmU2N2Rm",
  "test" : "dGVzdA=="
},
"kind" : "Secret",
"metadata" : {
  "creationTimestamp" : "2018-09-04T03:59:19Z",
  "name" : "secret-test",
  "namespace" : "namespace-test",
  "resourceVersion" : "5199399",
  "selfLink" : "/api/v1/namespaces/namespace-test/secrets/secret-test",
  "uid" : "e6170b6d-aff6-11e8-8f17-c81fbe371a17"
},
"type" : "Opaque"
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.11.6 更新 Secret

功能介绍

更新Secret中部分信息。

调用方法

请参见[如何调用API](#)。

URI

PATCH /api/v1/namespaces/{namespace}/secrets/{name}

表 5-3640 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Secret
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3641 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3642 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	目前支持三种类型的PATCH请求方法的操作，参考《 使用JSON合并patch更新Deployment 》。 <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-3643 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-3644 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3645 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3646 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3647 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

更新Secret中的数据为"test2" : "dGVzdA=="。

Content-Type: application/merge-patch+json

```
{
  "data": {
    "test2": "dGVzdA=="
  }
}
```


响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "data": {
    "key1": "MWYyZDFIMmU2N2Rm",
    "key2": "YWRtaW4=",
    "test2": "dGVzdA=="
  },
  "kind": "Secret",
  "metadata": {
    "creationTimestamp": "2018-09-04T03:59:19Z",
    "name": "secret-test",
    "namespace": "namespace-test",
    "resourceVersion": "5199399",
    "selfLink": "/api/v1/namespaces/namespace-test/secrets/secret-test",
    "uid": "e6170b6d-aff6-11e8-8f17-c81fbe371a17"
  },
  "type": "Opaque"
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.11.7 替换 Secret

功能介绍

替换Secret。

其中以下字段支持更新：

- metadata.labels
- metadata.annotations
- data

说明

说明：

- 当“type”的值为“Opaque”时，“data”的“key”和“value”都可以更新。
- 当“type”的值不为“Opaque”时，“data”的“value”可以更新。

调用方法

请参见[如何调用API](#)。

URI

PUT /api/v1/namespaces/{namespace}/secrets/{name}

表 5-3648 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Secret
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3649 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3650 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3651 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	是否必选	参数类型	描述
data	否	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	否	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by <code>ImmutableEphemeralVolumes</code> feature gate.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	否	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.

参数	是否必选	参数类型	描述
type	否	String	Used to facilitate programmatic handling of secret data.

表 5-3652 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	是否必选	参数类型	描述
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3653 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3654 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

响应参数

状态码： 200

表 5-3655 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '.'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3656 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3657 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3658 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

状态码： 201

表 5-3659 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
data	Map<String,String>	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string, representing the arbitrary (possibly non-string) data value here. Described in https://tools.ietf.org/html/rfc4648#section-4
immutable	Boolean	Immutable, if set to true, ensures that data stored in the Secret cannot be updated (only object metadata can be modified). If not set to true, the field can be modified at any time. Defaulted to nil. This is a beta field enabled by ImmutableEphemeralVolumes feature gate.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
stringData	Map<String,String>	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API.
type	String	Used to facilitate programmatic handling of secret data.

表 5-3660 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3661 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3662 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

请求示例

将已创建Secret中的key2替换为"test"。

```
{
  "apiVersion": "v1",
  "data": {
    "key1": "MWYyZDFIMmU2N2Rm",
    "test": "dGVzdA=="
  },
  "kind": "Secret",
  "metadata": {
```

```
"name" : "secret-test"  
},  
"type" : "Opaque"  
}
```

响应示例

状态码: 200

OK

```
{  
  "apiVersion" : "v1",  
  "data" : {  
    "key1" : "MWYyZDFlMmU2N2Rm",  
    "test" : "dGVzdA=="  
  },  
  "kind" : "Secret",  
  "metadata" : {  
    "creationTimestamp" : "2018-09-04T03:59:19Z",  
    "name" : "secret-test",  
    "namespace" : "namespace-test",  
    "resourceVersion" : "5199399",  
    "selfLink" : "/api/v1/namespaces/namespace-test/secrets/secret-test",  
    "uid" : "e6170b6d-aff6-11e8-8f17-c81fbe371a17"  
  },  
  "type" : "Opaque"  
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.12 Endpoint

5.12.1 查询指定 namespace 下的 Endpoints

功能介绍

查询Namespace下所有Endpoints。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/endpoints

表 5-3663 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3664 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3665 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3666 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Endpoints objects	List of endpoints.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3667 io.k8s.api.core.v1.Endpoints

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3668 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3669 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3670 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3671 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3672 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3673 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3674 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

表 5-3675 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2018-01-13T08:40:21Z",
      "name": "kubernetes",
      "namespace": "default",
      "resourceVersion": "49",
      "selfLink": "/api/v1/namespaces/default/endpoints/kubernetes",
      "uid": "64593b5d-f83d-11e7-9c3c-fa163eb8ad1a"
    },
    "subsets": [ {
      "addresses": [ {
        "ip": "192.168.0.64"
      } ],
      "ports": [ {
        "name": "https",
        "port": 5444,
        "protocol": "TCP"
      } ]
    } ]
  } ],
  "kind": "EndpointsList",
  "metadata": {
    "resourceVersion": "598704",
    "selfLink": "/api/v1/namespaces/default/endpoints"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.12.2 创建 Endpoint

功能介绍

创建Endpoint。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/endpoints

表 5-3676 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3677 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3678 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3679 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	否	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3680 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3681 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3682 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3683 io.k8s.api.core.v1.EndpointSubset

参数	是否必选	参数类型	描述
addresses	否	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	是否必选	参数类型	描述
notReadyAddresses	否	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	否	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3684 io.k8s.api.core.v1.EndpointAddress

参数	是否必选	参数类型	描述
hostname	否	String	The Hostname of this endpoint
ip	是	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	否	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	否	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3685 io.k8s.api.core.v1.ObjectReference

参数	是否必选	参数类型	描述
apiVersion	否	String	API version of the referent.

参数	是否必选	参数类型	描述
fieldPath	否	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	否	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	否	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	否	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency

参数	是否必选	参数类型	描述
uid	否	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3686 io.k8s.api.core.v1.EndpointPort

参数	是否必选	参数类型	描述
appProtocol	否	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Un-prefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	否	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	是	Integer	The port number of the endpoint.
protocol	否	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

响应参数

状态码： 200

表 5-3687 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3688 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3689 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3690 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3691 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3692 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3693 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3694 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

状态码： 201

表 5-3695 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3696 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3697 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3698 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3699 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3700 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3701 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3702 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

状态码： 202

表 5-3703 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3704 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3705 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3706 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3707 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3708 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3709 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3710 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

请求示例

创建Endpoint，端口名称为"service0"，端口号"80"，端口的IP协议为默认的TCP协议。

```
{
  "apiVersion": "v1",
  "kind": "Endpoints",
  "metadata": {
    "name": "test-endpoint",
    "namespace": "test-namespace"
  },
  "subsets": [ {
    "addresses": [ {
      "ip": "192.168.226.222",
      "targetRef": {
        "kind": "Pod",
        "name": "nginx-686bccd6f-ccsqb",
        "namespace": "test-namespace",
        "resourceVersion": "262055491",
        "uid": "ae417968-4f69-40da-9228-2434671e9c66"
      }
    }
  ]
},
  "ports": [ {
    "name": "service0",
    "port": 80,
    "protocol": "TCP"
  }
]
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Endpoints",
  "metadata": {
    "creationTimestamp": "2023-03-28T13:55:52Z",
    "name": "test-endpoint",
    "namespace": "test-namespace",
    "resourceVersion": "263538673",
    "selfLink": "/api/v1/namespaces/test-namespace/endpoints/test-endpoint",
    "uid": "3907aa64-c3f3-4de0-b938-9d23e8d12f94"
  },
  "subsets": [ {
    "addresses": [ {
```

```
"ip" : "192.168.226.222",
"targetRef" : {
  "kind" : "Pod",
  "name" : "nginx-686bccd6f-ccsqb",
  "namespace" : "test-namespace",
  "resourceVersion" : "262055491",
  "uid" : "ae417968-4f69-40da-9228-2434671e9c66"
},
}],
"ports" : [ {
  "name" : "service0",
  "port" : 80,
  "protocol" : "TCP"
} ]
} ]
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.12.3 删除 Endpoint

功能介绍

删除Endpoint。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/endpoints/{name}

表 5-3711 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Endpoints
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3712 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3713 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3714 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3715 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3716 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3717 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3718 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3719 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-3720 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3721 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3722 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3723 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "details": {
    "kind": "endpoints",
    "name": "test-endpoint",
    "uid": "3907aa64-c3f3-4de0-b938-9d23e8d12f94"
  },
  "kind": "Status",
```

```
"metadata": { },  
"status": "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.12.4 查询 Endpoint

功能介绍

查询Endpoint。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/endpoints/{name}

表 5-3724 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Endpoints
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3725 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3726 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3727 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3728 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3729 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3730 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3731 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3732 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3733 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3734 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Endpoints",
  "metadata": {
    "creationTimestamp": "2016-08-03T09:56:10Z",
    "name": "cluster-test",
    "namespace": "default",
    "resourceVersion": "18186",
    "selfLink": "/api/v1/namespaces/default/endpoints/cluster-test",
    "uid": "81b1503d-5960-11e6-b444-286ed488fafa"
  },
  "subsets": [
    {
      "addresses": [
        {
          "ip": "172.16.106.152"
        }
      ],
      "ports": [
        {
          "port": 1,
          "protocol": "TCP"
        }
      ]
    }
  ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.12.5 更新 Endpoint

功能介绍

更新Endpoint。

调用方法

请参见[如何调用API](#)。

URI

PATCH /api/v1/namespaces/{namespace}/endpoints/{name}

表 5-3735 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Endpoints
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3736 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3737 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-3738 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-3739 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3740 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3741 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3742 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3743 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3744 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3745 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3746 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

请求示例

更新Endpoint。

```
[{
  "op": "add",
  "path": "/metadata/annotations",
  "value": {
    "test-key": "test-value"
  }
}]
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Endpoints",
  "metadata": {
    "annotations": {
      "test-key": "test-value"
    }
  },
  "creationTimestamp": "2023-03-28T13:55:52Z",
  "name": "test-endpoint",
  "namespace": "test-namespace",
  "resourceVersion": "263561274",
  "selfLink": "/api/v1/namespaces/test-namespace/endpoints/test-endpoint",
  "uid": "3907aa64-c3f3-4de0-b938-9d23e8d12f94"
},
"subsets": [ {
  "addresses": [ {
    "ip": "192.168.226.222",
    "targetRef": {
      "kind": "Pod",
      "name": "nginx-686bccd6f-ccsqb",
      "namespace": "test-namespace",
      "resourceVersion": "262055491",
      "uid": "ae417968-4f69-40da-9228-2434671e9c66"
    }
  }
  ],
  "ports": [ {
    "name": "service0",
    "port": 80,
    "protocol": "TCP"
  }
  ]
}
]
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.12.6 替换 Endpoint

功能介绍

替换Endpoint。

调用方法

请参见[如何调用API](#)。

URI

PUT /api/v1/namespaces/{namespace}/endpoints/{name}

表 5-3747 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Endpoints
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3748 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3749 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3750 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	否	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3751 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3752 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3753 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3754 io.k8s.api.core.v1.EndpointSubset

参数	是否必选	参数类型	描述
addresses	否	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	是否必选	参数类型	描述
notReadyAddresses	否	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	否	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3755 io.k8s.api.core.v1.EndpointAddress

参数	是否必选	参数类型	描述
hostname	否	String	The Hostname of this endpoint
ip	是	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	否	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	否	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3756 io.k8s.api.core.v1.ObjectReference

参数	是否必选	参数类型	描述
apiVersion	否	String	API version of the referent.

参数	是否必选	参数类型	描述
fieldPath	否	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	否	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	否	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	否	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency

参数	是否必选	参数类型	描述
uid	否	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3757 io.k8s.api.core.v1.EndpointPort

参数	是否必选	参数类型	描述
appProtocol	否	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Un-prefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as mycompany.com/my-custom-protocol. This is a beta field that is guarded by the ServiceAppProtocol feature gate and enabled by default.
name	否	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	是	Integer	The port number of the endpoint.
protocol	否	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

响应参数

状态码： 200

表 5-3758 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3759 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3760 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3761 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3762 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3763 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3764 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3765 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

状态码： 201

表 5-3766 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
subsets	Array of io.k8s.api.core.v1.EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 5-3767 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3768 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3769 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3770 io.k8s.api.core.v1.EndpointSubset

参数	参数类型	描述
addresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of io.k8s.api.core.v1.EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of io.k8s.api.core.v1.EndpointPort objects	Port numbers available on the related IP addresses.

表 5-3771 io.k8s.api.core.v1.EndpointAddress

参数	参数类型	描述
hostname	String	The Hostname of this endpoint
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24). IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready.
nodeName	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	io.k8s.api.core.v1.ObjectReference object	Reference to object providing the endpoint.

表 5-3772 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3773 io.k8s.api.core.v1.EndpointPort

参数	参数类型	描述
appProtocol	String	The application protocol for this port. This field follows standard Kubernetes label syntax. Unprefixed names are reserved for IANA standard service names (as per RFC-6335 and http://www.iana.org/assignments/service-names). Non-standard protocols should use prefixed names such as <code>mycompany.com/my-custom-protocol</code> . This is a beta field that is guarded by the <code>ServiceAppProtocol</code> feature gate and enabled by default.

参数	参数类型	描述
name	String	The name of this port. This must match the 'name' field in the corresponding ServicePort. Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP, TCP, or SCTP. Default is TCP.

请求示例

替换Endpoint。

```
{
  "apiVersion": "v1",
  "kind": "Endpoints",
  "metadata": {
    "annotations": {
      "new-key": "new-value",
      "test-key": "test-value"
    },
    "creationTimestamp": "2023-03-28T13:55:52Z",
    "name": "test-endpoint",
    "namespace": "test-namespace",
    "resourceVersion": "263561274",
    "selfLink": "/api/v1/namespaces/test-namespace/endpoints/test-endpoint",
    "uid": "3907aa64-c3f3-4de0-b938-9d23e8d12f94"
  },
  "subsets": [ {
    "addresses": [ {
      "ip": "192.168.226.222",
      "targetRef": {
        "kind": "Pod",
        "name": "nginx-686bccd6f-ccsqb",
        "namespace": "test-namespace",
        "resourceVersion": "262055491",
        "uid": "ae417968-4f69-40da-9228-2434671e9c66"
      }
    }
  ]
},
  "ports": [ {
    "name": "service0",
    "port": 80,
    "protocol": "TCP"
  }
]
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "Endpoints",
  "metadata": {
    "annotations": {
      "new-key": "new-value",
```

```
    "test-key" : "test-value"
  },
  "creationTimestamp" : "2023-03-28T13:55:52Z",
  "name" : "test-endpoint",
  "namespace" : "test-namespace",
  "resourceVersion" : "263576210",
  "selfLink" : "/api/v1/namespaces/test-namespace/endpoints/test-endpoint",
  "uid" : "3907aa64-c3f3-4de0-b938-9d23e8d12f94"
},
"subsets" : [ {
  "addresses" : [ {
    "ip" : "192.168.226.222",
    "targetRef" : {
      "kind" : "Pod",
      "name" : "nginx-686bccd6f-ccsqb",
      "namespace" : "test-namespace",
      "resourceVersion" : "262055491",
      "uid" : "ae417968-4f69-40da-9228-2434671e9c66"
    }
  }
],
"ports" : [ {
  "name" : "service0",
  "port" : 80,
  "protocol" : "TCP"
}
]
}
]
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.13 ResourceQuota

5.13.1 查询指定 namespace 下的 resourcequotas

功能介绍

查询指定namespace下的resourcequotas（资源用量）。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/resourcequotas

表 5-3774 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3775 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3776 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3777 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.ResourceQuota objects	Items is a list of ResourceQuota objects. More info: https://kubernetes.io/docs/concepts/policy/resource-quotas/
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3778 io.k8s.api.core.v1.ResourceQuota

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ResourceQuotaSpec object	Spec defines the desired quota. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ResourceQuotaStatus object	Status defines the actual enforced quota and its current usage. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3779 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3780 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3781 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3782 io.k8s.api.core.v1.ResourceQuotaSpec

参数	参数类型	描述
hard	Map<String,String>	hard is the set of desired hard limits for each named resource. More info: https://kubernetes.io/docs/concepts/policy/resource-quotas/

参数	参数类型	描述
scopeSelector	io.k8s.api.core.v1.ScopeSelector object	scopeSelector is also a collection of filters like scopes that must match each object tracked by a quota but expressed using ScopeSelectorOperator in combination with possible values. For a resource to match, both scopes AND scopeSelector (if specified in spec), must be matched.
scopes	Array of strings	A collection of filters that must match each object tracked by a quota. If not specified, the quota matches all objects.

表 5-3783 io.k8s.api.core.v1.ScopeSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.ScopedResourceSelectorRequirement objects	A list of scope selector requirements by scope of the resources.

表 5-3784 io.k8s.api.core.v1.ScopedResourceSelectorRequirement

参数	参数类型	描述
operator	String	Represents a scope's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist.
scopeName	String	The name of the scope that the selector applies to.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3785 io.k8s.api.core.v1.ResourceQuotaStatus

参数	参数类型	描述
hard	Map<String,String>	Hard is the set of enforced hard limits for each named resource. More info: https://kubernetes.io/docs/concepts/policy/resource-quotas/
used	Map<String,String>	Used is the current observed total usage of the resource in the namespace.

表 5-3786 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2022-09-07T06:07:58Z",
      "name": "compute-resources",
      "namespace": "namespace-test",
      "resourceVersion": "42717140",
      "selfLink": "/api/v1/namespaces/namespace-test/resourcequotas/compute-resources",
      "uid": "df6e7f78-717d-48e2-bf41-afab0ba0a3de"
    },
    "spec": {
      "hard": {
        "configmaps": "480k",
        "count/cronjobs.batch": "960",
        "count/deployments.apps": "480",
        "count/ingresses.extensions": "480",
        "count/jobs.batch": "960",
        "count/networks.networking.cci.io": "72",
        "count/statefulsets.apps": "480",
        "limits.cpu": "100k",
        "limits.memory": "800000Gi",
        "persistentvolumeclaims": "960",
        "pods": "480k",
        "requests.nvidia.com/gpu-tesla-p4": "256",
        "requests.nvidia.com/gpu-tesla-p100-16GB": "256",
        "requests.nvidia.com/gpu-tesla-v100-16GB": "256",
        "requests.nvidia.com/gpu-tesla-v100-32GB": "256",
        "secrets": "480k",
        "services": "9600"
      }
    }
  }
]
```

```

    }
  },
  "status": {
    "hard": {
      "configmaps": "480k",
      "count/cronjobs.batch": "960",
      "count/deployments.apps": "480",
      "count/ingresses.extensions": "480",
      "count/jobs.batch": "960",
      "count/networks.networking.cci.io": "72",
      "count/statefulsets.apps": "480",
      "limits.cpu": "100k",
      "limits.memory": "800000Gi",
      "persistentvolumeclaims": "960",
      "pods": "480k",
      "requests.nvidia.com/gpu-tesla-p4": "256",
      "requests.nvidia.com/gpu-tesla-p100-16GB": "256",
      "requests.nvidia.com/gpu-tesla-v100-16GB": "256",
      "requests.nvidia.com/gpu-tesla-v100-32GB": "256",
      "secrets": "480k",
      "services": "9600"
    },
    "used": {
      "configmaps": "0",
      "count/cronjobs.batch": "0",
      "count/deployments.apps": "0",
      "count/ingresses.extensions": "0",
      "count/jobs.batch": "0",
      "count/networks.networking.cci.io": "1",
      "count/statefulsets.apps": "0",
      "limits.cpu": "0",
      "limits.memory": "0",
      "persistentvolumeclaims": "0",
      "pods": "0",
      "requests.nvidia.com/gpu-tesla-p4": "0",
      "requests.nvidia.com/gpu-tesla-p100-16GB": "0",
      "requests.nvidia.com/gpu-tesla-v100-16GB": "0",
      "requests.nvidia.com/gpu-tesla-v100-32GB": "0",
      "secrets": "2",
      "services": "1"
    }
  }
},
"kind": "ResourceQuotaList",
"metadata": {
  "resourceVersion": "42732540",
  "selfLink": "/api/v1/namespaces/namespace-test/resourcequotas"
}
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.13.2 查询指定的 resourcequota

功能介绍

查询指定的resourcequota（资源用量）。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/resourcequotas/{name}

表 5-3787 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ResourceQuota
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3788 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.

参数	是否必选	参数类型	描述
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3789 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3790 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.ResourceQuotaSpec object	Spec defines the desired quota. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.core.v1.ResourceQuotaStatus object	Status defines the actual enforced quota and its current usage. https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3791 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3792 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3793 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3794 io.k8s.api.core.v1.ResourceQuotaSpec

参数	参数类型	描述
hard	Map<String,String>	hard is the set of desired hard limits for each named resource. More info: https://kubernetes.io/docs/concepts/policy/resource-quotas/

参数	参数类型	描述
scopeSelector	io.k8s.api.core.v1.ScopeSelector object	scopeSelector is also a collection of filters like scopes that must match each object tracked by a quota but expressed using ScopeSelectorOperator in combination with possible values. For a resource to match, both scopes AND scopeSelector (if specified in spec), must be matched.
scopes	Array of strings	A collection of filters that must match each object tracked by a quota. If not specified, the quota matches all objects.

表 5-3795 io.k8s.api.core.v1.ScopeSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.ScopedResourceSelectorRequirement objects	A list of scope selector requirements by scope of the resources.

表 5-3796 io.k8s.api.core.v1.ScopedResourceSelectorRequirement

参数	参数类型	描述
operator	String	Represents a scope's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist.
scopeName	String	The name of the scope that the selector applies to.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3797 io.k8s.api.core.v1.ResourceQuotaStatus

参数	参数类型	描述
hard	Map<String,String>	Hard is the set of enforced hard limits for each named resource. More info: https://kubernetes.io/docs/concepts/policy/resource-quotas/
used	Map<String,String>	Used is the current observed total usage of the resource in the namespace.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "ResourceQuota",
  "metadata": {
    "creationTimestamp": "2022-09-07T06:07:58Z",
    "name": "compute-resources",
    "namespace": "namespace-test",
    "resourceVersion": "42717140",
    "selfLink": "/api/v1/namespaces/namespace-test/resourcequotas/compute-resources",
    "uid": "df6e7f78-717d-48e2-bf41-afab0ba0a3de"
  },
  "spec": {
    "hard": {
      "configmaps": "480k",
      "count/cronjobs.batch": "960",
      "count/deployments.apps": "480",
      "count/ingresses.extensions": "480",
      "count/jobs.batch": "960",
      "count/networks.networking.cci.io": "72",
      "count/statefulsets.apps": "480",
      "limits.cpu": "100k",
      "limits.memory": "800000Gi",
      "persistentvolumeclaims": "960",
      "pods": "480k",
      "requests.nvidia.com/gpu-tesla-p4": "256",
      "requests.nvidia.com/gpu-tesla-p100-16GB": "256",
      "requests.nvidia.com/gpu-tesla-v100-16GB": "256",
      "requests.nvidia.com/gpu-tesla-v100-32GB": "256",
      "secrets": "480k",
      "services": "9600"
    }
  },
  "status": {
    "hard": {
      "configmaps": "480k",
      "count/cronjobs.batch": "960",
      "count/deployments.apps": "480",
      "count/ingresses.extensions": "480",
      "count/jobs.batch": "960",
      "count/networks.networking.cci.io": "72",
      "count/statefulsets.apps": "480",

```

```
"limits.cpu" : "100k",
"limits.memory" : "800000Gi",
"persistentvolumeclaims" : "960",
"pods" : "480k",
"requests.nvidia.com/gpu-tesla-p4" : "256",
"requests.nvidia.com/gpu-tesla-p100-16GB" : "256",
"requests.nvidia.com/gpu-tesla-v100-16GB" : "256",
"requests.nvidia.com/gpu-tesla-v100-32GB" : "256",
"secrets" : "480k",
"services" : "9600"
},
"used" : {
  "configmaps" : "0",
  "count/cronjobs.batch" : "0",
  "count/deployments.apps" : "0",
  "count/ingresses.extensions" : "0",
  "count/jobs.batch" : "0",
  "count/networks.networking.cci.io" : "1",
  "count/statefulsets.apps" : "0",
  "limits.cpu" : "0",
  "limits.memory" : "0",
  "persistentvolumeclaims" : "0",
  "pods" : "0",
  "requests.nvidia.com/gpu-tesla-p4" : "0",
  "requests.nvidia.com/gpu-tesla-p100-16GB" : "0",
  "requests.nvidia.com/gpu-tesla-v100-16GB" : "0",
  "requests.nvidia.com/gpu-tesla-v100-32GB" : "0",
  "secrets" : "2",
  "services" : "1"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.14 CronJob

5.14.1 查询所有 namespace 下的 CronJobs

功能介绍

查询所有namespace下所有CronJob的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1beta1/cronjobs

表 5-3798 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 5-3799 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3800 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.batch.v1beta1.CronJob objects	items is the list of CronJobs.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-3801 io.k8s.api.batch.v1beta1.CronJob

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1beta1.CronJobSpec object	Specification of the desired behavior of a cron job, including the schedule. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1beta1.CronJobStatus object	Current status of a cron job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3802 io.k8s.api.batch.v1beta1.CronJobSpec

参数	参数类型	描述
concurrencyPolicy	String	Specifies how to treat concurrent executions of a Job. Valid values are: - "Allow" (default): allows CronJobs to run concurrently; - "Forbid": forbids concurrent runs, skipping next run if previous run hasn't finished yet; - "Replace": cancels currently running job and replaces it with a new one
failedJobsHistoryLimit	Integer	The number of failed finished jobs to retain. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1.
jobTemplate	io.k8s.api.batch.v1beta1.JobTemplateSpec object	Specifies the job that will be created when executing a CronJob.
schedule	String	The schedule in Cron format, see https://en.wikipedia.org/wiki/Cron .
startingDeadlineSeconds	Long	Optional deadline in seconds for starting the job if it misses scheduled time for any reason. Missed jobs executions will be counted as failed ones.
successfulJobsHistoryLimit	Integer	The number of successful finished jobs to retain. This is a pointer to distinguish between explicit zero and not specified. Defaults to 3.

参数	参数类型	描述
suspend	Boolean	This flag tells the controller to suspend subsequent executions, it does not apply to already started executions. Defaults to false.

表 5-3803 io.k8s.api.batch.v1beta1.JobTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata of the jobs created from this template. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of the job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3804 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((.spec.completions - .status.successful) < .spec.parallelism)$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-3805 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3806 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-3807 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-3808 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-3809 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-3810 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-3811 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-3812 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-3813 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-3814 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-3815 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-3816 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-3817 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-3818 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-3819 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-3820 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-3821 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-3822 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-3823 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-3824 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-3825 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-3826 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-3827 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-3828 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-3829 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-3830 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-3831 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-3832 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-3833 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-3834 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-3835 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-3836 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-3837 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-3838 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-3839 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-3840 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-3841 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain '%'. Must not contain '*'. Must not contain '?'. Must not contain '['. Must not contain ']'. Must not contain '&'. Must not contain '='. Must not contain '". Must not contain ''.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-3842 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-3843 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-3844 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-3845 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-3846 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-3847 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-3848 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-3849 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-3850 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-3851 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-3852 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-3853 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-3854 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-3855 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-3856 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-3857 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-3858 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-3859 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-3860 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-3861 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-3862 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-3863 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-3864 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3865 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3866 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3867 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-3868 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-3869 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-3870 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-3871 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-3872 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-3873 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-3874 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-3875 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-3876 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-3877 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-3878 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-3879 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-3880 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-3881 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-3882 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-3883 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-3884 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-3885 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-3886 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-3887 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-3888 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-3889 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-3890 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-3891 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-3892 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-3893 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-3894 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-3895 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-3896 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-3897 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-3898 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-3899 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-3900 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-3901 io.k8s.api.batch.v1beta1.CronJobStatus

参数	参数类型	描述
active	Array of io.k8s.api.core.v1.ObjectReference objects	A list of pointers to currently running jobs.
lastScheduleTime	String	Information when was the last time the job was successfully scheduled.

表 5-3902 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.

参数	参数类型	描述
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3903 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1beta1",
  "items": [ {
    "metadata": {
      "annotations": {
        "description": ""
      }
    },
    "creationTimestamp": "2022-09-06T08:56:37Z",
```

```
"name": "cronjob-test",
"namespace": "namespace-test",
"resourceVersion": "615089077",
"selfLink": "/apis/batch/v1beta1/namespaces/namespace-test/cronjobs/cronjob-test",
"uid": "808543b6-696f-49ef-a8c9-30ab6b66c979"
},
"spec": {
  "concurrencyPolicy": "Forbid",
  "failedJobsHistoryLimit": 1,
  "jobTemplate": {
    "metadata": {
      "creationTimestamp": null
    },
    "spec": {
      "template": {
        "metadata": {
          "annotations": {
            "cri.cci.io/container-type": "secure-container",
            "log.stdoutcollection.kubernetes.io": "{\"collectionContainers\": [\"container-0\"]}"
          },
          "creationTimestamp": null
        },
        "spec": {
          "containers": [ {
            "image": "redis",
            "imagePullPolicy": "IfNotPresent",
            "lifecycle": { },
            "name": "container-0",
            "resources": {
              "limits": {
                "cpu": "500m",
                "memory": "1Gi"
              },
              "requests": {
                "cpu": "500m",
                "memory": "1Gi"
              }
            },
            "terminationMessagePath": "/dev/termination-log",
            "terminationMessagePolicy": "File"
          } ],
          "dnsPolicy": "ClusterFirst",
          "imagePullSecrets": [ {
            "name": "imagepull-secret"
          } ],
          "restartPolicy": "Never",
          "schedulerName": "default-scheduler",
          "securityContext": { },
          "terminationGracePeriodSeconds": 30
        }
      }
    },
    "schedule": "0 */1 * * *",
    "successfulJobsHistoryLimit": 3,
    "suspend": false
  },
  "status": { }
}],
"kind": "CronJobList",
"metadata": {
  "resourceVersion": "41656236",
  "selfLink": ""
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15 API groups

5.15.1 查询 APIVersions

功能介绍

get available API versions

调用方法

请参见[如何调用API](#)。

URI

GET /api

请求参数

表 5-3904 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3905 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.

参数	参数类型	描述
versions	Array of strings	versions are the api versions that are available.

表 5-3906 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "kind": "APIVersions",
  "serverAddressByClientCIDRs": [ {
    "clientCIDR": "0.0.0.0/0",
    "serverAddress": "https://cci.cn-east-3.myhuaweicloud.com"
  } ],
  "versions": [ "v1" ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable

状态码	描述
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.2 查询所有 v1 版本的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1

请求参数

表 5-3907 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3908 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3909 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.

参数	参数类型	描述
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码： 200

OK

```
{
  "groupVersion": "v1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "ConfigMap",
    "name": "configmaps",
    "namespaced": true,
    "shortNames": [ "cm" ],
    "singularName": "",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }, {
    "kind": "Endpoints",
    "name": "endpoints",
    "namespaced": true,
    "shortNames": [ "ep" ],
    "singularName": "",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }
]
```

```
}, {
  "kind": "Event",
  "name": "events",
  "namespaced": true,
  "shortNames": [ "ev" ],
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "kind": "Namespace",
  "name": "namespaces",
  "namespaced": false,
  "shortNames": [ "ns" ],
  "singularName": "",
  "verbs": [ "create", "delete", "get", "list", "patch", "update", "watch" ]
}, {
  "kind": "PersistentVolumeClaim",
  "name": "persistentvolumeclaims",
  "namespaced": true,
  "shortNames": [ "pvc" ],
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "categories": [ "all" ],
  "kind": "Pod",
  "name": "pods",
  "namespaced": true,
  "shortNames": [ "po" ],
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "kind": "Pod",
  "name": "pods/exec",
  "namespaced": true,
  "singularName": "",
  "verbs": [ ]
}, {
  "kind": "Pod",
  "name": "pods/log",
  "namespaced": true,
  "singularName": "",
  "verbs": [ "get" ]
}, {
  "kind": "Pod",
  "name": "pods/status",
  "namespaced": true,
  "singularName": "",
  "verbs": [ "get", "patch", "update" ]
}, {
  "kind": "ResourceQuota",
  "name": "resourcequotas",
  "namespaced": true,
  "shortNames": [ "quota" ],
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "kind": "Secret",
  "name": "secrets",
  "namespaced": true,
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "categories": [ "all" ],
  "kind": "Service",
  "name": "services",
  "namespaced": true,
  "shortNames": [ "svc" ],
  "singularName": "",
  "verbs": [ "create", "delete", "get", "list", "patch", "update", "watch" ]
}, {
```

```
"kind" : "Service",  
"name" : "services/status",  
"namespaced" : true,  
"singularName" : "",  
"verbs" : [ "get", "patch", "update" ]  
}]  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.3 查询 APIGroupList

功能介绍

get available API versions

调用方法

请参见[如何调用API](#)。

URI

GET /apis

请求参数

表 5-3910 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3911 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groups	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIGroup objects	groups is a list of APIGroup.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3912 io.k8s.apimachinery.pkg.apis.meta.v1.APIGroup

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3913 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3914 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groups": [ {
    "name": "extensions",
    "preferredVersion": {
      "groupVersion": "extensions/v1beta1",
      "version": "v1beta1"
    },
    "versions": [ {
      "groupVersion": "extensions/v1beta1",
      "version": "v1beta1"
    } ]
  }, {
    "name": "apps",
    "preferredVersion": {
      "groupVersion": "apps/v1",
      "version": "v1"
    },
    "versions": [ {
      "groupVersion": "apps/v1",
      "version": "v1"
    } ],
    {
      "groupVersion": "apps/v1beta1",
      "version": "v1beta1"
    }
  } ]
}
```

```
}, {
  "name": "batch",
  "preferredVersion": {
    "groupVersion": "batch/v1",
    "version": "v1"
  },
  "versions": [ {
    "groupVersion": "batch/v1",
    "version": "v1"
  } ],
  {
    "groupVersion": "batch/v1beta1",
    "version": "v1beta1"
  } ]
}, {
  "name": "rbac.authorization.k8s.io",
  "preferredVersion": {
    "groupVersion": "rbac.authorization.k8s.io/v1",
    "version": "v1"
  },
  "versions": [ {
    "groupVersion": "rbac.authorization.k8s.io/v1",
    "version": "v1"
  } ]
}, {
  "name": "networking.cci.io",
  "preferredVersion": {
    "groupVersion": "networking.cci.io/v1beta1",
    "version": "v1beta1"
  },
  "versions": [ {
    "groupVersion": "networking.cci.io/v1beta1",
    "version": "v1beta1"
  } ]
}, {
  "name": "image.cci.io",
  "preferredVersion": {
    "groupVersion": "image.cci.io/v1alpha1",
    "version": "v1alpha1"
  },
  "versions": [ {
    "groupVersion": "image.cci.io/v1alpha1",
    "version": "v1alpha1"
  } ]
}, {
  "name": "kubeflow.org",
  "preferredVersion": {
    "groupVersion": "kubeflow.org/v1",
    "version": "v1"
  },
  "versions": [ {
    "groupVersion": "kubeflow.org/v1",
    "version": "v1"
  } ]
}, {
  "name": "batch.volcano.sh",
  "preferredVersion": {
    "groupVersion": "batch.volcano.sh/v1alpha1",
    "version": "v1alpha1"
  },
  "versions": [ {
    "groupVersion": "batch.volcano.sh/v1alpha1",
    "version": "v1alpha1"
  } ]
}, {
  "name": "metrics.k8s.io",
  "preferredVersion": {
    "groupVersion": "metrics.k8s.io/v1beta1",
    "version": "v1beta1"
  },
  "versions": [ {
    "groupVersion": "metrics.k8s.io/v1beta1",
    "version": "v1beta1"
  } ]
},
```



```
"versions": [{
  "groupVersion": "metrics.k8s.io/v1beta1",
  "version": "v1beta1"
}]
}],
"kind": "APIGroupList"
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.4 查询 APIGroup /apis/apps

功能介绍

查询APIGroup /apis/apps

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps

请求参数

表 5-3915 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3916 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.

参数	参数类型	描述
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3917 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3918 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "apps",
  "preferredVersion": {
    "groupVersion": "apps/v1",
    "version": "v1"
  },
  "versions": [ {
    "groupVersion": "apps/v1",
    "version": "v1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.5 查询所有 apps/v1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1

请求参数

表 5-3919 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3920 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3921 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.

参数	参数类型	描述
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "apps/v1",
  "kind": "APIResourceList",
  "resources": [ {
    "categories": [ "all" ],
    "kind": "Deployment",
    "name": "deployments",
    "namespaced": true,
    "shortNames": [ "deploy" ],
    "singularName": "",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }, {
    "group": "autoscaling",
    "kind": "Scale",
    "name": "deployments/scale",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "patch", "update" ],
    "version": "v1"
  }, {
    "kind": "Deployment",
    "name": "deployments/status",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "patch", "update" ]
  } ]
}
```

```
}, {
  "categories": [ "all" ],
  "kind": "ReplicaSet",
  "name": "replicasets",
  "namespaced": true,
  "shortNames": [ "rs" ],
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "categories": [ "all" ],
  "kind": "StatefulSet",
  "name": "statefulsets",
  "namespaced": true,
  "shortNames": [ "sts" ],
  "singularName": "",
  "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
}, {
  "kind": "StatefulSet",
  "name": "statefulsets/status",
  "namespaced": true,
  "singularName": "",
  "verbs": [ "get", "patch", "update" ]
}]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.6 查询 APIGroup /apis/batch

功能介绍

查询APIGroup /apis/batch

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch

请求参数

表 5-3922 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3923 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.

参数	参数类型	描述
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3924 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3925 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.

参数	参数类型	描述
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "batch",
  "preferredVersion": {
    "groupVersion": "batch/v1",
    "version": "v1"
  },
  "versions": [ {
    "groupVersion": "batch/v1",
    "version": "v1"
  }, {
    "groupVersion": "batch/v1beta1",
    "version": "v1beta1"
  }, {
    "groupVersion": "batch/v2alpha1",
    "version": "v2alpha1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.7 查询 APIGroup /apis/batch.volcano.sh

功能介绍

查询APIGroup /apis/batch.volcano.sh

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch.volcano.sh

请求参数

表 5-3926 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3927 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3928 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3929 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "batch.volcano.sh",
  "preferredVersion": {
    "groupVersion": "batch.volcano.sh/v1alpha1",
    "version": "v1alpha1"
  },
  "versions": [ {
    "groupVersion": "batch.volcano.sh/v1alpha1",
    "version": "v1alpha1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest

状态码	描述
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.8 查询所有 batch.volcano.sh/v1alpha1 的 API

功能介绍

查询所有batch.volcano.sh/v1alpha1的API

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch.volcano.sh/v1alpha1

请求参数

表 5-3930 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3931 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3932 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.

参数	参数类型	描述
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "batch.volcano.sh/v1alpha1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "Job",
    "name": "jobs",
    "namespaced": true,
    "shortNames": [ "vcjob", "vj" ],
    "singularName": "job",
    "storageVersionHash": "DbMX/QqtM30=",
    "verbs": [ "delete", "deletecollection", "get", "list", "patch", "create", "update", "watch" ]
  }, {
```

```
"kind" : "Job",  
"name" : "jobs/status",  
"namespaced" : true,  
"singularName" : "",  
"verbs" : [ "get", "patch", "update" ]  
}]  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.9 查询所有 batch/v1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1

请求参数

表 5-3933 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3934 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3935 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "batch/v1",
  "kind": "APIResourceList",
  "resources": [ {
    "categories": [ "all" ],
    "kind": "Job",
    "name": "jobs",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }, {
    "kind": "Job",
    "name": "jobs/status",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "patch", "update" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.10 查询所有 batch/v1beta1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1beta1

请求参数

表 5-3936 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3937 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3938 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.

参数	参数类型	描述
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "batch/v1beta1",
  "kind": "APIResourceList",
  "resources": [ {
    "categories": [ "all" ],
    "kind": "CronJob",
    "name": "cronjobs",
    "namespaced": true,
    "shortNames": [ "cj" ],
    "singularName": "",
    "storageVersionHash": "h/JlFAZkyy=",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest

状态码	描述
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.11 查询 APIGroup /apis/crd.yangtse.cni

功能介绍

查询APIGroup /apis/crd.yangtse.cni

调用方法

请参见[如何调用API](#)。

URI

GET /apis/crd.yangtse.cni

请求参数

表 5-3939 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3940 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.

参数	参数类型	描述
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3941 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3942 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "crd.yangtse.cni",
  "preferredVersion": {
    "groupVersion": "crd.yangtse.cni/v1",
    "version": "v1"
  },
}
```

```
"versions": [ {  
  "groupVersion": "crd.yangtse.cni/v1",  
  "version": "v1"  
} ]  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.12 查询所有 crd.yangtse.cni/v1 的 API

功能介绍

查询所有 crd.yangtse.cni/v1 的 API

调用方法

请参见[如何调用API](#)。

URI

GET /apis/crd.yangtse.cni/v1

请求参数

表 5-3943 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3944 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3945 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "crd.yangtse.cni/v1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "EIPPool",
    "name": "eippools",
    "namespaced": true,
    "singularName": "eippool",
    "storageVersionHash": "vG508QhjuOY=",
    "verbs": [ "delete", "deletecollection", "get", "list", "patch", "create", "update", "watch" ]
  }, {
    "kind": "EIPPool",
    "name": "eippools/status",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "patch", "update" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.13 查询 APIGroup /apis/extensions

功能介绍

查询APIGroup /apis/extensions

调用方法

请参见[如何调用API](#)。

URI

GET /apis/extensions

请求参数

表 5-3946 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3947 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.

参数	参数类型	描述
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-IP header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3948 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3949 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.

参数	参数类型	描述
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "extensions",
  "preferredVersion": {
    "groupVersion": "extensions/v1beta1",
    "version": "v1beta1"
  },
  "versions": [ {
    "groupVersion": "extensions/v1beta1",
    "version": "v1beta1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError

状态码	描述
503	ServiceUnavailable
504	ServerTimeout

5.15.14 查询所有 extensions/v1beta1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/extensions/v1beta1

请求参数

表 5-3950 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3951 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
groupVersion	String	groupVersion is the group and version this APIResourceList is for.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3952 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.

参数	参数类型	描述
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "groupVersion": "extensions/v1beta1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "Ingress",
    "name": "ingresses",
    "namespaced": true,
    "shortNames": [ "ing" ],
    "singularName": "",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }, {
    "kind": "Ingress",
    "name": "ingresses/status",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "patch", "update" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.15 查询 APIGroup /apis/metrics.k8s.io

功能介绍

查询APIGroup /apis/metrics.k8s.io

调用方法

请参见[如何调用API](#)。

URI

GET /apis/metrics.k8s.io

请求参数

表 5-3953 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3954 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.

参数	参数类型	描述
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-IP header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3955 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3956 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "metrics.k8s.io",
  "preferredVersion": {
    "groupVersion": "metrics.k8s.io/v1beta1",
    "version": "v1beta1"
  },
  "versions": [ {
    "groupVersion": "metrics.k8s.io/v1beta1",
    "version": "v1beta1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.16 查询所有 metrics.k8s.io/v1beta1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/metrics.k8s.io/v1beta1

请求参数

表 5-3957 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3958 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3959 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.

参数	参数类型	描述
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "metrics.k8s.io/v1beta1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "NodeMetrics",
    "name": "nodes",
    "namespaced": false,
    "singularName": "",
    "verbs": [ "get", "list" ]
  }, {
    "kind": "PodMetrics",
    "name": "pods",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "list" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.17 查询 APIGroup /apis/networking.cci.io

功能介绍

查询APIGroup /apis/networking.cci.io

调用方法

请参见[如何调用API](#)。

URI

GET /apis/networking.cci.io

请求参数

表 5-3960 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3961 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.

参数	参数类型	描述
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3962 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3963 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "networking.cci.io",
  "preferredVersion": {
    "groupVersion": "networking.cci.io/v1beta1",
    "version": "v1beta1"
  },
  "versions": [ {
    "groupVersion": "networking.cci.io/v1beta1",
    "version": "v1beta1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.18 查询所有 networking.cci.io/v1beta1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/networking.cci.io/v1beta1

请求参数

表 5-3964 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3965 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3966 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.

参数	参数类型	描述
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "networking.cci.io/v1beta1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "Network",
    "name": "networks",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }, {
    "kind": "Network",
    "name": "networks/status",
    "namespaced": true,
    "singularName": "",
    "verbs": [ "get", "patch", "update" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.19 查询 APIGroup /apis/rbac.authorization.k8s.io

功能介绍

get information of a group

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io

请求参数

表 5-3967 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3968 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	name is the name of the group.
preferredVersion	io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery object	preferredVersion is the version preferred by the API server, which probably is the storage version.

参数	参数类型	描述
serverAddressByClientCIDRs	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR objects	a map of client CIDR to server address that is serving this group. This is to help clients reach servers in the most network-efficient way possible. Clients can use the appropriate server address as per the CIDR that they match. In case of multiple matches, clients should use the longest matching CIDR. The server returns only those CIDRs that it thinks that the client can match. For example: the master will return an internal IP CIDR only, if the client reaches the server using an internal IP. Server looks at X-Forwarded-For header or X-Real-Ip header or request.RemoteAddr (in that order) to get the client IP.
versions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery objects	versions are the versions supported in this group.

表 5-3969 io.k8s.apimachinery.pkg.apis.meta.v1.GroupVersionForDiscovery

参数	参数类型	描述
groupVersion	String	groupVersion specifies the API group and version in the form "group/version"
version	String	version specifies the version in the form of "version". This is to save the clients the trouble of splitting the GroupVersion.

表 5-3970 io.k8s.apimachinery.pkg.apis.meta.v1.ServerAddressByClientCIDR

参数	参数类型	描述
clientCIDR	String	The CIDR with which clients can match their IP to figure out the server address that they should use.
serverAddress	String	Address of this server, suitable for a client that matches the above CIDR. This can be a hostname, hostname:port, IP or IP:port.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "APIGroup",
  "name": "rbac.authorization.k8s.io",
  "preferredVersion": {
    "groupVersion": "rbac.authorization.k8s.io/v1",
    "version": "v1"
  },
  "versions": [ {
    "groupVersion": "rbac.authorization.k8s.io/v1",
    "version": "v1"
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.15.20 查询所有 rbac.authorization.k8s.io/v1 的 API

功能介绍

get available resources

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io/v1

请求参数

表 5-3971 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3972 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
groupVersion	String	groupVersion is the group and version this APIResourceList is for.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
resources	Array of io.k8s.apimachinery.pkg.apis.meta.v1.APIResource objects	resources contains the name of the resources and if they are namespaced.

表 5-3973 io.k8s.apimachinery.pkg.apis.meta.v1.APIResource

参数	参数类型	描述
categories	Array of strings	categories is a list of the grouped resources this resource belongs to (e.g. 'all')
group	String	group is the preferred group of the resource. Empty implies the group of the containing resource list. For subresources, this may have a different value, for example: Scale".
kind	String	kind is the kind for the resource (e.g. 'Foo' is the kind for a resource 'foo')
name	String	name is the plural name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not.
shortNames	Array of strings	shortNames is a list of suggested short names of the resource.
singularName	String	singularName is the singular name of the resource. This allows clients to handle plural and singular opaquely. The singularName is more correct for reporting status on a single item and both singular and plural are allowed from the kubectl CLI interface.

参数	参数类型	描述
storageVersionHash	String	The hash value of the storage version, the version this resource is converted to when written to the data store. Value must be treated as opaque by clients. Only equality comparison on the value is valid. This is an alpha feature and may change or be removed in the future. The field is populated by the apiserver only if the StorageVersionHash feature gate is enabled. This field will remain optional even if it graduates.
verbs	Array of strings	verbs is a list of supported kube verbs (this includes get, list, watch, create, update, patch, delete, deletecollection, and proxy)
version	String	version is the preferred version of the resource. Empty implies the version of the containing resource list For subresources, this may have a different value, for example: v1 (while inside a v1beta1 version of the core resource's group)".

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "groupVersion": "rbac.authorization.k8s.io/v1",
  "kind": "APIResourceList",
  "resources": [ {
    "kind": "ClusterRole",
    "name": "clusterroles",
    "namespaced": false,
    "singularName": "",
    "storageVersionHash": "bYE5ZWDrJ44=",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  }, {
    "kind": "RoleBinding",
    "name": "rolebindings",
    "namespaced": true,
    "singularName": "",
    "storageVersionHash": "eGsCzGH6b1g=",
    "verbs": [ "create", "delete", "deletecollection", "get", "list", "patch", "update", "watch" ]
  } ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.16 Event

5.16.1 查询指定 namespace 下的 Events

功能介绍

查询Namespace下所有Event的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/events

表 5-3974 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3975 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3976 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-3977 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.Event objects	List of events
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-3978 io.k8s.api.core.v1.Event

参数	参数类型	描述
action	String	What action was taken/failed regarding to the Regarding object.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
count	Integer	The number of times this event has occurred.
eventTime	String	Time when this Event was first observed.

参数	参数类型	描述
firstTimestamp	String	The time at which the event was first recorded. (Time of server receipt is in TypeMeta.)
involvedObject	io.k8s.api.core.v1.ObjectReference object	The object that this event is about.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
lastTimestamp	String	The time at which the most recent occurrence of this event was recorded.
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
reason	String	This should be a short, machine understandable string that gives the reason for the transition into the object's current status.
related	io.k8s.api.core.v1.ObjectReference object	Optional secondary object for more complex actions.
reportingComponent	String	Name of the controller that emitted this Event, e.g. <i>kubernetes.io/kubelet</i> .
reportingInstance	String	ID of the controller instance, e.g. <i>kubelet-xyzf</i> .
series	io.k8s.api.core.v1.EventSeries object	Data about the Event series this event represents or nil if it's a singleton Event.
source	io.k8s.api.core.v1.EventSource object	The component reporting this event. Should be a short machine understandable string.
type	String	Type of this event (Normal, Warning), new types could be added in the future

表 5-3979 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-3980 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-3981 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3982 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-3983 io.k8s.api.core.v1.EventSeries

参数	参数类型	描述
count	Integer	Number of occurrences in this series up to the last heartbeat time
lastObservedTime	String	Time of the last occurrence observed

表 5-3984 io.k8s.api.core.v1.EventSource

参数	参数类型	描述
component	String	Component from which the event is generated.
host	String	Node name on which the event is generated.

表 5-3985 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "count": 1,
    "eventTime": null,
    "firstTimestamp": "2018-09-03T12:58:07Z",
    "involvedObject": {
      "apiVersion": "v1",
      "kind": "Pod",
      "name": "deployment-test-57f7cff77c-5x5tw",
      "namespace": "namespace-test",
      "resourceVersion": "5036865",
      "uid": "010fec39-af79-11e8-8f17-c81fbe371a17"
    },
    },
  "lastTimestamp": "2018-09-03T12:58:07Z",
  "message": "Successfully assigned deployment-test-57f7cff77c-5x5tw to\nc0dd6256-195a-e811-90a2-10c17294fcbc",
  "metadata": {
    "creationTimestamp": "2018-09-03T12:58:07Z",
    "name": "deployment-test-57f7cff77c-5x5tw.1550e534d2d8a5ef",
    "namespace": "namespace-test",
    "resourceVersion": "760533",
    "selfLink": "/api/v1/namespaces/namespace-test/events/deployment-test-57f7cff77c-5x5tw.1550e534d2d8a5ef",
    "uid": "0122b5b2-af79-11e8-8f17-c81fbe371a17"
  },
  "reason": "Scheduled",
  "reportingComponent": "",
  "reportingInstance": "",
  "source": {
```

```
    "component" : "default-scheduler"
  },
  "type" : "Normal"
}, {
  "count" : 2,
  "eventTime" : null,
  "firstTimestamp" : "2018-09-03T12:58:07Z",
  "involvedObject" : {
    "apiVersion" : "v1",
    "kind" : "Pod",
    "name" : "deployment-test-57f7cff77c-5x5tw",
    "namespace" : "namespace-test",
    "resourceVersion" : "5036870",
    "uid" : "010fec39-af79-11e8-8f17-c81fbe371a17"
  },
  "lastTimestamp" : "2018-09-03T12:58:12Z",
  "message" : "Successfully mounted volumes for pod \"deployment-test-57f7cff77c-5x5tw_namespace-test(010fec39-af79-11e8-8f17-c81fbe371a17)\"",
  "metadata" : {
    "creationTimestamp" : "2018-09-03T12:58:07Z",
    "name" : "deployment-test-57f7cff77c-5x5tw.1550e534d3105acd",
    "namespace" : "namespace-test",
    "resourceVersion" : "760542",
    "selfLink" : "/api/v1/namespaces/namespace-test/events/deployment-test-57f7cff77c-5x5tw.1550e534d3105acd",
    "uid" : "0123faea-af79-11e8-8f17-c81fbe371a17"
  },
  "reason" : "SuccessfulMountVolume",
  "reportingComponent" : "",
  "reportingInstance" : "",
  "source" : {
    "component" : "kubelet",
    "host" : "c0dd6256-195a-e811-90a2-10c17294fcbc"
  },
  "type" : "Normal"
}, {
  "count" : 1,
  "eventTime" : null,
  "firstTimestamp" : "2018-09-03T12:58:09Z",
  "involvedObject" : {
    "apiVersion" : "v1",
    "fieldPath" : "spec.containers{container-0}",
    "kind" : "Pod",
    "name" : "deployment-test-57f7cff77c-5x5tw",
    "namespace" : "namespace-test",
    "resourceVersion" : "5036870",
    "uid" : "010fec39-af79-11e8-8f17-c81fbe371a17"
  },
  "lastTimestamp" : "2018-09-03T12:58:09Z",
  "message" : "pulling image \"*:*:20202/cci/redis:V1\"",
  "metadata" : {
    "creationTimestamp" : "2018-09-03T12:58:09Z",
    "name" : "deployment-test-57f7cff77c-5x5tw.1550e5354a5915e0",
    "namespace" : "namespace-test",
    "resourceVersion" : "760536",
    "selfLink" : "/api/v1/namespaces/namespace-test/events/deployment-test-57f7cff77c-5x5tw.1550e5354a5915e0",
    "uid" : "025554a4-af79-11e8-8f17-c81fbe371a17"
  },
  "reason" : "Pulling",
  "reportingComponent" : "",
  "reportingInstance" : "",
  "source" : {
    "component" : "kubelet",
    "host" : "c0dd6256-195a-e811-90a2-10c17294fcbc"
  },
  "type" : "Normal"
}
}],
"kind" : "EventList",
```

```
"metadata": {  
  "resourceVersion": "764693",  
  "selfLink": "/api/v1/namespaces/namespace-test/events"  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.16.2 删除 Event

功能介绍

删除Event。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/events/{name}

表 5-3986 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Event
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-3987 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-3988 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-3989 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-3990 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-3991 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3992 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3993 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3994 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-3995 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-3996 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-3997 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-3998 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "details": {
    "kind": "events",
    "name": "cci-deployment-cml002.15c0dbdd5715a8da",
  }
}
```

```
"uid" : "68db3125-ce19-11e9-8d71-d0efc1b3bb6b"  
},  
"kind" : "Status",  
"metadata" : { },  
"status" : "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.16.3 查询 Event

功能介绍

查询Event详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/events/{name}

表 5-3999 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Event
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4000 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4001 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4002 响应 Body 参数

参数	参数类型	描述
action	String	What action was taken/failed regarding to the Regarding object.

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
count	Integer	The number of times this event has occurred.
eventTime	String	Time when this Event was first observed.
firstTimestamp	String	The time at which the event was first recorded. (Time of server receipt is in TypeMeta.)
involvedObject	io.k8s.api.core.v1.ObjectReference object	The object that this event is about.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
lastTimestamp	String	The time at which the most recent occurrence of this event was recorded.
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
reason	String	This should be a short, machine understandable string that gives the reason for the transition into the object's current status.
related	io.k8s.api.core.v1.ObjectReference object	Optional secondary object for more complex actions.
reportingComponent	String	Name of the controller that emitted this Event, e.g. <i>kubernetes.io/kubelet</i> .
reportingInstance	String	ID of the controller instance, e.g. <i>kubelet-xyzf</i> .

参数	参数类型	描述
series	io.k8s.api.core.v1.EventSeries object	Data about the Event series this event represents or nil if it's a singleton Event.
source	io.k8s.api.core.v1.EventSource object	The component reporting this event. Should be a short machine understandable string.
type	String	Type of this event (Normal, Warning), new types could be added in the future

表 5-4003 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4004 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4005 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4006 io.k8s.api.core.v1.ObjectReference

参数	参数类型	描述
apiVersion	String	API version of the referent.

参数	参数类型	描述
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
namespace	String	Namespace of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/namespaces/
resourceVersion	String	Specific resourceVersion to which this reference is made, if any. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
uid	String	UID of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#uids

表 5-4007 io.k8s.api.core.v1.EventSeries

参数	参数类型	描述
count	Integer	Number of occurrences in this series up to the last heartbeat time
lastObservedTime	String	Time of the last occurrence observed

表 5-4008 io.k8s.api.core.v1.EventSource

参数	参数类型	描述
component	String	Component from which the event is generated.
host	String	Node name on which the event is generated.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "count": 1,
  "eventTime": null,
  "firstTimestamp": "2018-09-03T12:58:07Z",
  "involvedObject": {
    "apiVersion": "v1",
    "kind": "Pod",
    "name": "deployment-test-57f7cff77c-5x5tw",
    "namespace": "namespace-test",
    "resourceVersion": "5036865",
    "uid": "010fec39-af79-11e8-8f17-c81fbe371a17"
  },
  "kind": "Event",
  "lastTimestamp": "2018-09-03T12:58:07Z",
  "message": "Successfully assigned deployment-test-57f7cff77c-5x5tw to c0dd6256-195a-
e811-90a2-10c17294fcbc",
  "metadata": {
    "creationTimestamp": "2018-09-03T12:58:07Z",
    "name": "deployment-test-57f7cff77c-5x5tw.1550e534d2d8a5ef",
    "namespace": "namespace-test",
    "resourceVersion": "760533",
    "selfLink": "/api/v1/namespaces/namespace-test/events/deployment-
test-57f7cff77c-5x5tw.1550e534d2d8a5ef",
    "uid": "0122b5b2-af79-11e8-8f17-c81fbe371a17"
  },
  "reason": "Scheduled",
  "reportingComponent": "",
  "reportingInstance": "",
  "source": {
    "component": "default-scheduler"
  },
  "type": "Normal"
}
```

状态码

状态码	描述
200	OK
400	BadRequest

状态码	描述
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.17 PersistentVolumeClaim

5.17.1 查询指定 namespace 下的 PersistentVolumeClaims

功能介绍

查询Namespace下的所有PersistentVolumeClaim。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/persistentvolumeclaims

表 5-4009 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4010 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4011 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4012 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	A list of persistent volume claims. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-4013 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4014 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4015 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4016 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4017 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4018 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4019 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4020 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4021 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4022 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4023 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-4024 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "items": [ {
    "metadata": {
      "annotations": {
        "pv.kubernetes.io/bind-completed": "yes",
        "pv.kubernetes.io/bound-by-controller": "yes",
        "volume.beta.kubernetes.io/storage-class": "sata",
        "volume.beta.kubernetes.io/storage-provisioner": "flexvolume-huawei.com/fuxivol"
      },
      "creationTimestamp": "2018-09-07T07:17:38Z",
      "finalizers": [ "kubernetes.io/pvc-protection" ],
      "name": "pvc-test",
      "namespace": "test-namespace",
      "resourceVersion": "5915795",
      "selfLink": "/api/v1/namespaces/test-namespace/persistentvolumeclaims/pvc-test",
      "uid": "19a355cc-b26e-11e8-b205-c88d83be759f"
    },
    "spec": {
      "accessModes": [ "ReadWriteMany" ],
      "resources": {
        "requests": {
          "storage": "10Gi"
        }
      }
    },
    "volumeName": "pvc-19a355cc-b26e-11e8-b205-c88d83be759f"
  } ],
  "status": {
    "accessModes": [ "ReadWriteMany" ],
    "capacity": {
```

```
    "storage": "10Gi"
  },
  "phase": "Bound"
}
}],
"kind": "PersistentVolumeClaimList",
"metadata": {
  "resourceVersion": "6456754",
  "selfLink": "/api/v1/namespaces/namespace-test/persistentvolumeclaims"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.17.2 创建 PersistentVolumeClaim

功能介绍

创建PersistentVolumeClaim，主要适用于动态创建存储的场景，即存储资源未创建时，创建PVC会根据请求内容创建一个存储资源。

当前支持创建EVS（云硬盘卷，块存储）和SFS（文件存储卷），使用时spec.storageClassName参数的取值如下：

- sata：普通I/O云硬盘卷
- sas：高I/O云硬盘卷
- ssd：超高I/O云硬盘卷

- csi-disk-gpssd: 通用型SSD云硬盘卷
- csi-disk-gpssd2: 通用型SSD V2云硬盘卷
- nfs-rw: 标准文件协议类型文件存储卷
- csi-sfs: SFS 3.0容量型弹性文件服务
- csi-sfsturbo-hpc: HPC型弹性文件服务

若不指定spec.storageClassName, 可在metadata.annotations中指定volume.beta.kubernetes.io/storage-class, 取值含义与spec.storageClassName 相同, 这两个参数选择其中一个即可。

若需要创建加密类型存储卷, 对于云硬盘存储卷需要在metadata.annotations中增加paas.storage.io/cryptKeyId字段; 对于文件存储卷需要增加paas.storage.io/cryptKeyId、paas.storage.io/cryptAlias和paas.storage.io/cryptDomainId字段。

若需要创建通用型SSD V2云硬盘卷, 卷需要在metadata.annotations中增加everest.io/disk-iops 和 everest.io/disk-throughput 字段, IOPS和吞吐量大小范围可见EVS文档。

调用方法

请参见[如何调用API](#)。

URI

POST /api/v1/namespaces/{namespace}/persistentvolumeclaims

表 5-4025 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4026 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4027 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4028 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	是否必选	参数类型	描述
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	否	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4029 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	是否必选	参数类型	描述
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4030 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4031 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4032 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4033 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-4034 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4035 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4036 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4037 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	否	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	否	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	否	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4038 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	是	String	status is the status of the condition.
type	是	String	type is the type of the condition.

响应参数

状态码： 200

表 5-4039 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4040 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4041 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4042 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4043 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4044 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4045 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4046 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4047 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4048 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4049 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

状态码： 201

表 5-4050 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4051 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4052 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4053 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4054 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4055 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4056 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4057 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4058 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4059 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4060 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

状态码： 202

表 5-4061 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4062 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4063 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4064 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4065 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4066 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4067 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4068 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4069 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4070 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4071 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

请求示例

- 创建一个大小为10G的加密云硬盘，类型为sata

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "paas.storage.io/cryptKeyId": "ee9b610c-e356-11e9-aadc-d0efc1b3bb6b",
      "volume.beta.kubernetes.io/storage-class": "sata"
    },
    "name": "pvc-test",
    "namespace": "test-namespace"
  },
  "spec": {
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    }
  }
}
```

- 创建一个大小为10G的加密文件存储卷

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "paas.storage.io/cryptAlias": "sfs/default",
      "paas.storage.io/cryptDomainId": "d6912480-c3d6-4e9e-8c70-38afeea434c3",
      "paas.storage.io/cryptKeyId": "ee9b610c-e356-11e9-aadc-d0efc1b3bb6b",
      "volume.beta.kubernetes.io/storage-class": "nfs-rw"
    },
    "name": "pvc-test",
    "namespace": "test-namespace"
  },
  "spec": {
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "paas.storage.io/cryptKeyId": "ee9b610c-e356-11e9-aadc-d0efc1b3bb6b",
      "pv.kubernetes.io/bind-completed": "yes",
      "pv.kubernetes.io/bound-by-controller": "yes",
      "volume.beta.kubernetes.io/storage-provisioner": "flexvolume-huawei.com/fuxivol"
    },
    "creationTimestamp": "2018-11-24T07:48:35Z",
    "finalizers": [ "kubernetes.io/pvc-protection" ],
    "labels": {
      "app": "evs"
    },
    "name": "pvc-test",
    "namespace": "test-namespace",
    "resourceVersion": "28156856",
    "selfLink": "/api/v1/namespaces/ns-test/persistentvolumeclaims/pvc-test",
    "uid": "58d15f3e-efbd-11e8-8950-501d934409f3"
  },
  "spec": {
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    },
    "storageClassName": "sata",
    "volumeName": "pvc-58d15f3e-efbd-11e8-8950-501d934409f3"
  },
  "status": {
    "accessModes": [ "ReadWriteMany" ],
    "capacity": {
      "storage": "10Gi"
    },
    "phase": "Bound"
  }
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.17.3 删除 PersistentVolumeClaim

功能介绍

删除PersistentVolumeClaim。

调用方法

请参见[如何调用API](#)。

URI

DELETE /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}

表 5-4072 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the PersistentVolumeClaim
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4073 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4074 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4075 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-4076 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-4077 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4078 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4079 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4080 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4081 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4082 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4083 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4084 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4085 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4086 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4087 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

状态码： 202

表 5-4088 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4089 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4090 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4091 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4092 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4093 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4094 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4095 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4096 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4097 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4098 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

请求示例

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions"
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "volume.beta.kubernetes.io/storage-class": "sata",
      "volume.beta.kubernetes.io/storage-provisioner": "flexvolume-huawei.com/fuxivol"
    },
    "creationTimestamp": "2018-09-06T07:27:47Z",
    "deletionGracePeriodSeconds": 0,
    "deletionTimestamp": "2018-09-06T07:51:51Z",
    "finalizers": [ "kubernetes.io/pvc-protection" ],
    "name": "pvc-test",
    "namespace": "namespace-test",
    "resourceVersion": "5669396",
    "selfLink": "/api/v1/namespaces/namespace-test/persistentvolumeclaims/pvc-test",
    "uid": "5a9ab4c1-b1a6-11e8-8b84-c88d83be759f"
  },
  "spec": {
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    }
  },
  "status": {
    "phase": "Pending"
  }
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.17.4 查询 PersistentVolumeClaim

功能介绍

查询PersistentVolumeClaim。

调用方法

请参见[如何调用API](#)。

URI

GET /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}

表 5-4099 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the PersistentVolumeClaim

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4100 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4101 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4102 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4103 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4104 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4105 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4106 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4107 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4108 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4109 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4110 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4111 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4112 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "pv.kubernetes.io/bind-completed": "yes",
      "pv.kubernetes.io/bound-by-controller": "yes",
      "volume.beta.kubernetes.io/storage-class": "sata",
      "volume.beta.kubernetes.io/storage-provisioner": "flexvolume-huawei.com/fuxivol"
    },
    "creationTimestamp": "2018-09-07T07:17:38Z",
    "finalizers": [ "kubernetes.io/pvc-protection" ],
    "name": "pvc-test",
    "namespace": "test-namespace",
    "resourceVersion": "5915795",
    "selfLink": "/api/v1/namespaces/test-namespace/persistentvolumeclaims/pvc-test",
    "uid": "19a355cc-b26e-11e8-b205-c88d83be759f"
  },
  "spec": {
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    }
  },
  "volumeName": "pvc-19a355cc-b26e-11e8-b205-c88d83be759f"
},
  "status": {
    "accessModes": [ "ReadWriteMany" ],
    "capacity": {
      "storage": "10Gi"
    },
    "phase": "Bound"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.17.5 替换 PersistentVolumeClaims

功能介绍

替换指定PersistentVolumeClaims。

其中以下字段支持更新：

- metadata.labels
- metadata.annotations
- metadata.generateName 其余部分不支持更新。

调用方法

请参见[如何调用API](#)。

URI

PUT /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}

表 5-4113 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the PersistentVolumeClaim
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4114 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4115 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4116 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

参数	是否必选	参数类型	描述
status	否	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4117 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4118 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4119 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4120 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4121 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-4122 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4123 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4124 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4125 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	否	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	否	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	否	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4126 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	是	String	status is the status of the condition.
type	是	String	type is the type of the condition.

响应参数

状态码： 200

表 5-4127 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4128 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4129 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4130 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4131 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4132 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4133 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4134 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4135 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4136 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4137 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

状态码： 201

表 5-4138 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4139 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4140 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4141 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4142 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4143 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4144 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4145 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4146 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4147 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4148 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

请求示例

更新PersistentVolumeClaims的metadata.labels

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "everest.io/disk-volume-type": "SATA",
      "pv.kubernetes.io/bind-completed": "yes",
      "pv.kubernetes.io/bound-by-controller": "yes",
      "volume.beta.kubernetes.io/storage-provisioner": "everest-csi-provisioner"
    },
    "creationTimestamp": "2023-06-25T02:59:26Z",
    "finalizers": [ "kubernetes.io/pvc-protection" ],
    "labels": {
      "failure-domain.beta.kubernetes.io/region": "cn-north-7",
      "failure-domain.beta.kubernetes.io/zone": "cn-north-7a"
    },
    "name": "pvc-test",
    "namespace": "test-namespace",
    "resourceVersion": "343613181",
    "selfLink": "/api/v1/namespaces/test-namespace/persistentvolumeclaims/pvc-test",
    "uid": "bb99f33d-a5da-46b4-90a4-b0c6b3eedabf"
  },
  "spec": {
    "accessModes": [ "ReadWriteMany" ],
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    },
    "storageClassName": "sata",
    "volumeMode": "Filesystem",
    "volumeName": "pvc-bb99f33d-a5da-46b4-90a4-b0c6b3eedabf"
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
```

```

"everest.io/disk-volume-type" : "SATA",
"pv.kubernetes.io/bind-completed" : "yes",
"pv.kubernetes.io/bound-by-controller" : "yes",
"volume.beta.kubernetes.io/storage-provisioner" : "everest-csi-provisioner"
},
"creationTimestamp" : "2023-06-25T02:59:26Z",
"finalizers" : [ "kubernetes.io/pvc-protection" ],
"labels" : {
  "failure-domain.beta.kubernetes.io/region" : "cn-north-7",
  "failure-domain.beta.kubernetes.io/zone" : "cn-north-7a"
},
"name" : "pvc-test",
"namespace" : "test-namespace",
"resourceVersion" : "343622839",
"selfLink" : "/api/v1/namespaces/test-namespace/persistentvolumeclaims/pvc-test",
"uid" : "bb99f33d-a5da-46b4-90a4-b0c6b3eedabf"
},
"spec" : {
  "accessModes" : [ "ReadWriteMany" ],
  "resources" : {
    "requests" : {
      "storage" : "10Gi"
    }
  },
  "storageClassName" : "sata",
  "volumeMode" : "Filesystem",
  "volumeName" : "pvc-bb99f33d-a5da-46b4-90a4-b0c6b3eedabf"
},
"status" : {
  "accessModes" : [ "ReadWriteMany" ],
  "capacity" : {
    "storage" : "10Gi"
  },
  "phase" : "Bound"
}
}

```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests

状态码	描述
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18 RoleBinding

5.18.1 获取指定 namespace 下 RoleBinding 列表

功能介绍

This API is used to list or watch objects of kind RoleBinding

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings

表 5-4149 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4150 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4151 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4152 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.rbac.v1.RoleBinding objects	Items is a list of RoleBindings
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard object's metadata.

表 5-4153 io.k8s.api.rbac.v1.RoleBinding

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4154 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4155 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4156 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4157 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4158 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

表 5-4159 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2018-11-27T03:29:52Z",
      "name": "read-pods",
      "namespace": "default",
      "resourceVersion": "16611",
```

```
"selfLink" : "/apis/rbac.authorization.k8s.io/v1/namespaces/default/rolebindings/secret-reader",
"uid" : "b3d1a49a-f1f4-11e8-b449-fa163ec24e06"
},
"roleRef" : {
"apiGroup" : "rbac.authorization.k8s.io",
"kind" : "Role",
"name" : "pod-reader"
},
"subjects" : [ {
"apiGroup" : "rbac.authorization.k8s.io",
"kind" : "User",
"name" : "jane"
} ]
} ],
"kind" : "RoleBindingList",
"metadata" : {
"resourceVersion" : "16611",
"selfLink" : "/apis/rbac.authorization.k8s.io/v1/namespaces/default/rolebindings"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18.2 创建 RoleBinding

功能介绍

This API is used to create a RoleBinding

调用方法

请参见[如何调用API](#)。

URI

POST /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings

表 5-4160 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4161 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4162 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4163 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.

参数	是否必选	参数类型	描述
roleRef	是	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	否	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4164 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4165 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4166 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4167 io.k8s.api.rbac.v1.RoleRef

参数	是否必选	参数类型	描述
apiGroup	是	String	APIGroup is the group for the resource being referenced
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-4168 io.k8s.api.rbac.v1.Subject

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	是	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	是	String	Name of the object being referenced.
namespace	否	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

响应参数

状态码： 200

表 5-4169 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4170 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4171 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4172 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4173 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4174 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

状态码： 201

表 5-4175 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4176 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4177 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4178 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4179 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4180 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

状态码： 202

表 5-4181 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4182 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4183 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4184 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4185 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4186 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

请求示例

创建RoleBinding，将"rbac-test"命名空间中的用户组绑定到名称为"view"的ClusterRole下。

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "kind": "RoleBinding",
  "metadata": {
    "name": "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "namespace": "rbac-test"
  },
  "roleRef": {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "ClusterRole",
    "name": "view"
  },
  "subjects": [ {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "User",
    "name": "07b82a44a680d5661f01c00b448f8f50"
  } ]
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "kind": "RoleBinding",
  "metadata": {
    "creationTimestamp": "2020-04-07T08:25:46Z",
    "name": "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "namespace": "rbac-test",
    "resourceVersion": "230511279",
    "selfLink": "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "uid": "6163c216-78a9-11ea-bcc5-340a9837e2a7"
  },
  "roleRef": {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "ClusterRole",
    "name": "view"
  },
  "subjects": [ {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "User",
    "name": "07b82a44a680d5661f01c00b448f8f50"
  } ]
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18.3 删除指定的 RoleBinding

功能介绍

This API is used to delete a RoleBinding

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}

表 5-4187 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the RoleBinding
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4188 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4189 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4190 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-4191 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-4192 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4193 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4194 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-4195 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-4196 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4197 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4198 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-4199 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "v1",
  "details": {
    "group": "rbac.authorization.k8s.io",
    "kind": "rolebindings",
    "name": "read-pods",
    "uid": "b3d1a49a-f1f4-11e8-b449-fa163ec24e06"
  },
}
```

```
"kind": "Status",  
"metadata": {},  
"status": "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18.4 获取指定的 RoleBinding

功能介绍

This API is used to read the specified RoleBinding

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}

表 5-4200 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the RoleBinding
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4201 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4202 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4203 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4204 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4205 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4206 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4207 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4208 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognized the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "kind": "RoleBinding",
  "metadata": {
    "creationTimestamp": "2020-04-07T08:25:46Z",
    "name": "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "namespace": "rbac-test",
    "resourceVersion": "230511279",
    "selfLink": "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "uid": "6163c216-78a9-11ea-bcc5-340a9837e2a7"
  },
  "roleRef": {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "ClusterRole",
    "name": "view"
  },
  "subjects": [ {
```

```
"apiGroup" : "rbac.authorization.k8s.io",  
"kind" : "User",  
"name" : "07b82a44a680d5661f01c00b448f8f50"  
}]  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18.5 更新指定的 RoleBinding

功能介绍

This API is used to partially update the specified RoleBinding

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}

表 5-4209 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the RoleBinding
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4210 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4211 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	目前支持三种类型的PATCH请求方法的操作，参考《 使用JSON合并patch更新Deployment 》。 <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-4212 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-4213 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4214 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4215 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4216 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4217 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4218 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognized the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

请求示例

更新指定的RoleBinding。

```
{
  "roleRef": {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "ClusterRole",
    "name": "view"
  },
  "subjects": [ {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "User",
    "name": "07b82a44a680d5661f01c00b448f8f50"
  }, {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "Group",
    "name": "07b8387fd080d2963f92c00bb9a7341e"
  } ]
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
```



```
"kind": "RoleBinding",
"metadata": {
  "creationTimestamp": "2020-04-07T08:25:46Z",
  "name": "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
  "namespace": "rbac-test",
  "resourceVersion": "230608015",
  "selfLink": "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
  "uid": "6163c216-78a9-11ea-bcc5-340a9837e2a7"
},
"roleRef": {
  "apiGroup": "rbac.authorization.k8s.io",
  "kind": "ClusterRole",
  "name": "view"
},
"subjects": [ {
  "apiGroup": "rbac.authorization.k8s.io",
  "kind": "User",
  "name": "07b82a44a680d5661f01c00b448f8f50"
}, {
  "apiGroup": "rbac.authorization.k8s.io",
  "kind": "Group",
  "name": "07b8387fd080d2963f92c00bb9a7341e"
} ]
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18.6 替换指定的 RoleBinding

功能介绍

This API is used to replace the specified RoleBinding

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}

表 5-4219 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the RoleBinding
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4220 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4221 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4222 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.

参数	是否必选	参数类型	描述
roleRef	是	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	否	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4223 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4224 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4225 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4226 io.k8s.api.rbac.v1.RoleRef

参数	是否必选	参数类型	描述
apiGroup	是	String	APIGroup is the group for the resource being referenced
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-4227 io.k8s.api.rbac.v1.Subject

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	是	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	是	String	Name of the object being referenced.
namespace	否	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

响应参数

状态码： 200

表 5-4228 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4229 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4230 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4231 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4232 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4233 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

状态码： 201

表 5-4234 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4235 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4236 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4237 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4238 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4239 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

请求示例

替换指定的RoleBinding。

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "kind": "RoleBinding",
  "metadata": {
    "creationTimestamp": "2020-04-07T08:25:46Z",
    "name": "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "namespace": "rbac-test",
    "resourceVersion": "230608015",
    "selfLink": "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "uid": "6163c216-78a9-11ea-bcc5-340a9837e2a7"
  },
  "roleRef": {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "ClusterRole",
    "name": "view"
  },
  "subjects": [ {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "User",
    "name": "07b82a44a680d5661f01c00b448f8f50"
  } ]
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "kind": "RoleBinding",
  "metadata": {
    "creationTimestamp": "2020-04-07T08:25:46Z",
    "name": "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "namespace": "rbac-test",
    "resourceVersion": "230609819",
    "selfLink": "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
    "uid": "6163c216-78a9-11ea-bcc5-340a9837e2a7"
  },
  "roleRef": {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "ClusterRole",
    "name": "view"
  },
  "subjects": [ {
    "apiGroup": "rbac.authorization.k8s.io",
    "kind": "User",
    "name": "07b82a44a680d5661f01c00b448f8f50"
  } ]
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.18.7 获取 RoleBinding 列表

功能介绍

This API is used to list or watch objects of kind RoleBinding

调用方法

请参见[如何调用API](#)。

URI

GET /apis/rbac.authorization.k8s.io/v1/rolebindings

表 5-4240 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 5-4241 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4242 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.rbac.v1.RoleBinding objects	Items is a list of RoleBindings
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard object's metadata.

表 5-4243 io.k8s.api.rbac.v1.RoleBinding

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata.
roleRef	io.k8s.api.rbac.v1.RoleRef object	RoleRef can reference a Role in the current namespace or a ClusterRole in the global namespace. If the RoleRef cannot be resolved, the Authorizer must return an error.
subjects	Array of io.k8s.api.rbac.v1.Subject objects	Subjects holds references to the objects the role applies to.

表 5-4244 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

参数	参数类型	描述
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4245 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4246 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4247 io.k8s.api.rbac.v1.RoleRef

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced
kind	String	Kind is the type of resource being referenced

参数	参数类型	描述
name	String	Name is the name of resource being referenced

表 5-4248 io.k8s.api.rbac.v1.Subject

参数	参数类型	描述
apiGroup	String	APIGroup holds the API group of the referenced subject. Defaults to "" for ServiceAccount subjects. Defaults to "rbac.authorization.k8s.io" for User and Group subjects.
kind	String	Kind of object being referenced. Values defined by this API group are "User", "Group", and "ServiceAccount". If the Authorizer does not recognize the kind value, the Authorizer should report an error.
name	String	Name of the object being referenced.
namespace	String	Namespace of the referenced object. If the object kind is non-namespace, such as "User" or "Group", and this value is not empty the Authorizer should report an error.

表 5-4249 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "rbac.authorization.k8s.io/v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2020-04-07T08:24:24Z",
      "name": "clusterrole_cluster-admin_User_456b80acdfd1471e8a7fbc0019825124",
      "namespace": "rbac-test",
      "resourceVersion": "230509842",
```

```
"selfLink" : "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_cluster-admin_User_456b80acdf1471e8a7fbc0019825124",
"uid" : "303fcfde-78a9-11ea-83a1-340a9837e413"
},
"roleRef" : {
"apiGroup" : "rbac.authorization.k8s.io",
"kind" : "ClusterRole",
"name" : "cluster-admin"
},
"subjects" : [ {
"apiGroup" : "rbac.authorization.k8s.io",
"kind" : "User",
"name" : "456b80acdf1471e8a7fbc0019825124"
} ]
}, {
"metadata" : {
"creationTimestamp" : "2020-04-07T08:25:46Z",
"name" : "clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
"namespace" : "rbac-test",
"resourceVersion" : "230511279",
"selfLink" : "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings/clusterrole_view_User_07b82a44a680d5661f01c00b448f8f50",
"uid" : "6163c216-78a9-11ea-bcc5-340a9837e2a7"
},
"roleRef" : {
"apiGroup" : "rbac.authorization.k8s.io",
"kind" : "ClusterRole",
"name" : "view"
},
"subjects" : [ {
"apiGroup" : "rbac.authorization.k8s.io",
"kind" : "User",
"name" : "07b82a44a680d5661f01c00b448f8f50"
} ]
},
"kind" : "RoleBindingList",
"metadata" : {
"resourceVersion" : "230535192",
"selfLink" : "/apis/rbac.authorization.k8s.io/v1/namespaces/rbac-test/rolebindings"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid

状态码	描述
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19 StatefulSet

5.19.1 删除指定 namespace 下的 StatefulSets

功能介绍

删除Namespace下所有StatefulSet。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/apps/v1/namespaces/{namespace}/statefulsets

表 5-4250 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4251 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4252 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4253 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-4254 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-4255 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4256 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4257 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-4258 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

- 只删除StatefulSet（对应Pod不删除）

```
{
  "Kind": "DeleteOptions",
  "apiVersion": "v1",
  "propagationPolicy": "Orphan"
}
```

- 前台级联删除（按照Pod->StatefulSet的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Foreground"
}
```

- 后台级联删除（按照StatefulSet->Pod的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Background"
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "apps/v1",
  "items": null,
  "kind": "StatefulSetList",
  "metadata": {
    "resourceVersion": "5257614",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.2 查询指定 namespace 下的 StatefulSets

功能介绍

查询Namespace下所有StatefulSet的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/statefulsets

表 5-4259 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4260 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4261 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4262 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.apps.v1.StatefulSet objects	StatefulSet represents a set of pods with consistent identities. Identities are defined as: <ul style="list-style-type: none">• Network: A single stable DNS and hostname.• Storage: As many VolumeClaims as requested. The StatefulSet guarantees that a given network identity will always map to the same storage identity.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	ListMeta describes metadata that synthetic resources must have, including lists and various status objects. A resource may have only one of {ObjectMeta, ListMeta}.

表 5-4263 io.k8s.api.apps.v1.StatefulSet

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4264 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.

参数	参数类型	描述
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4265 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4266 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4267 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4268 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4269 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4270 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4271 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4272 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4273 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4274 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4275 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4276 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4277 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4278 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-4279 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4280 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-4281 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4282 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4283 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4284 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4285 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4286 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4287 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4288 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-4289 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-4290 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4291 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4292 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4293 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4294 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-4295 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4296 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4297 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4298 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-4299 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4300 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4301 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4302 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4303 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4304 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4305 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4306 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-4307 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4308 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-4309 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4310 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4311 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4312 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4313 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4314 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4315 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-4316 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4317 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4318 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4319 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4320 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4321 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4322 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4323 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4324 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4325 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4326 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4327 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4328 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-4329 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4330 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4331 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4332 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-4333 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4334 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4335 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-4336 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-4337 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4338 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4339 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4340 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4341 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4342 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-4343 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-4344 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4345 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-4346 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-4347 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-4348 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-4349 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-4350 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-4351 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-4352 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-4353 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-4354 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-4355 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4356 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4357 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4358 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4359 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4360 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4361 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4362 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4363 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4364 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4365 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-4366 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-4367 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

表 5-4368 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2018-09-04T07:13:00Z",
      "generation": 1,
      "labels": {
        "app": "statefulset-test"
      }
    }
  }
]
```

```
},
"name": "statefulset-test",
"namespace": "namespace-test",
"resourceVersion": "5209881",
"selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets/statefulset-test",
"uid": "f4a35f35-b011-11e8-b6ef-f898ef6c78b4"
},
"spec": {
  "podManagementPolicy": "OrderedReady",
  "replicas": 3,
  "revisionHistoryLimit": 10,
  "selector": {
    "matchLabels": {
      "app": "statefulset-test"
    }
  }
},
"serviceName": "",
"template": {
  "metadata": {
    "annotations": {
      "cri.cci.io/container-type": "secure-container"
    },
    "creationTimestamp": null,
    "labels": {
      "app": "statefulset-test"
    }
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "IfNotPresent",
      "name": "container-0",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      },
      "terminationMessagePath": "/dev/termination-log",
      "terminationMessagePolicy": "File"
    } ],
    "dnsPolicy": "ClusterFirst",
    "imagePullSecrets": [ {
      "name": "imagepull-secret"
    } ],
    "restartPolicy": "Always",
    "schedulerName": "default-scheduler",
    "securityContext": { }
  }
},
"updateStrategy": {
  "type": "OnDelete"
},
"status": {
  "collisionCount": 0,
  "currentReplicas": 3,
  "currentRevision": "statefulset-test-f986b645b",
  "observedGeneration": 1,
  "readyReplicas": 2,
  "replicas": 3,
  "updateRevision": "statefulset-test-f986b645b"
}
}],
"kind": "StatefulSetList",
```

```
"metadata": {  
  "resourceVersion": "5215730",  
  "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets"  
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.3 创建 StatefulSet

功能介绍

创建StatefulSet。

调用方法

请参见[如何调用API](#)。

URI

POST /apis/apps/v1/namespaces/{namespace}/statefulsets

表 5-4369 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4370 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4371 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4372 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	否	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	否	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4373 io.k8s.api.apps.v1.StatefulSetSpec

参数	是否必选	参数类型	描述
podManagementPolicy	否	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	否	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	否	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	是	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	是否必选	参数类型	描述
serviceName	是	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	是	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	否	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	否	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4374 io.k8s.api.core.v1.PodTemplateSpec

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4375 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	是否必选	参数类型	描述
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	是否必选	参数类型	描述
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod

参数	是否必选	参数类型	描述
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	是否必选	参数类型	描述
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.

参数	是否必选	参数类型	描述
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy

参数	是否必选	参数类型	描述
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	否	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	是否必选	参数类型	描述
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	否	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.

参数	是否必选	参数类型	描述
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4376 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	是否必选	参数类型	描述
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4377 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4378 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4379 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4380 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4381 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4382 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-4383 io.k8s.api.core.v1.PodAntiAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-4384 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4385 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4386 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4387 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.
value	否	String	value is the value of the option

表 5-4388 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	是否必选	参数类型	描述
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4389 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-4390 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4391 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4392 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4393 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4394 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4395 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-4396 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4397 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-4398 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-4399 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4400 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4401 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4402 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4403 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-4404 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4405 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4406 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4407 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-4408 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4409 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4410 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain '!'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4411 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4412 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4413 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4414 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4415 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-4416 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4417 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-4418 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
gmsaCredentialSpecName	否	String	<code>GMSACredentialSpecName</code> is the name of the GMSA credential spec to use.
runAsUserName	否	String	The <code>UserName</code> in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in <code>PodSecurityContext</code> . If set in both <code>SecurityContext</code> and <code>PodSecurityContext</code> , the value specified in <code>SecurityContext</code> takes precedence.

表 5-4419 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are <code>NoSchedule</code> , <code>PreferNoSchedule</code> and <code>NoExecute</code> .
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be <code>Exists</code> ; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4420 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4421 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4422 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4423 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4424 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-4425 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4426 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4427 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4428 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4429 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4430 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4431 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4432 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4433 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4434 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4435 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4436 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4437 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-4438 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4439 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4440 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.

参数	是否必选	参数类型	描述
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4441 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or</p>

参数	是否必选	参数类型	描述
			<p>suffix) is as large as possible. The sign will be omitted unless the number is negative. Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-4442 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4443 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4444 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-4445 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-4446 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4447 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4448 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4449 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4450 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4451 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-4452 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-4453 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-4454 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-4455 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-4456 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-4457 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-4458 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-4459 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-4460 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-4461 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmk

表 5-4462 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	是否必选	参数类型	描述
rollingUpdate	否	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	否	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-4463 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	是否必选	参数类型	描述
partition	否	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-4464 io.k8s.api.core.v1.PersistentVolumeClaim

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	否	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4465 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	否	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4466 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4467 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4468 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4469 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-4470 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4471 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4472 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4473 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	否	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	否	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	否	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4474 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	是	String	status is the status of the condition.
type	是	String	type is the type of the condition.

表 5-4475 io.k8s.api.apps.v1.StatefulSetStatus

参数	是否必选	参数类型	描述
collisionCount	否	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	否	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	否	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	否	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).

参数	是否必选	参数类型	描述
observedGeneration	否	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	否	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	是	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	否	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)
updatedReplicas	否	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-4476 io.k8s.api.apps.v1.StatefulSetCondition

参数	是否必选	参数类型	描述
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.
message	否	String	A human readable message indicating details about the transition.
reason	否	String	The reason for the condition's last transition.
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of statefulset condition.

响应参数

状态码： 200

表 5-4477 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4478 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.

参数	参数类型	描述
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4479 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4480 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4481 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4482 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4483 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4484 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4485 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4486 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4487 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4488 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4489 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4490 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4491 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4492 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-4493 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4494 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-4495 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4496 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4497 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4498 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4499 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4500 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4501 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4502 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-4503 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-4504 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4505 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4506 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4507 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4508 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-4509 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4510 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4511 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4512 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-4513 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4514 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4515 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4516 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4517 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4518 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4519 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4520 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-4521 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4522 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-4523 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4524 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4525 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4526 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4527 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4528 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4529 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-4530 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4531 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4532 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4533 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4534 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4535 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4536 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4537 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4538 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4539 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4540 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4541 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4542 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-4543 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4544 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4545 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4546 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-4547 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4548 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4549 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-4550 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-4551 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4552 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4553 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4554 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4555 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4556 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-4557 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-4558 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4559 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-4560 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-4561 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-4562 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-4563 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-4564 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-4565 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-4566 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-4567 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-4568 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-4569 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4570 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4571 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4572 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4573 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4574 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4575 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4576 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4577 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4578 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4579 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-4580 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-4581 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

状态码： 201

表 5-4582 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.

参数	参数类型	描述
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4583 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	参数类型	描述
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4584 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4585 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.

参数	参数类型	描述
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	参数类型	描述
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

参数	参数类型	描述
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://github.com/kubernetes/enhancements/blob/master/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	参数类型	描述
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4586 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4587 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4588 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4589 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4590 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4591 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4592 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4593 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4594 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4595 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4596 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4597 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-4598 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4599 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-4600 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4601 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4602 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4603 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4604 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4605 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4606 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4607 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-4608 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-4609 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4610 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4611 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4612 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4613 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-4614 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4615 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4616 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4617 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-4618 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4619 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4620 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '..'. Must not contain '/'. Must not contain '%'. Must not contain '*'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4621 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4622 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4623 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4624 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4625 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-4626 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4627 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-4628 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4629 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4630 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4631 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	<p>FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.</p>
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	<p>Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running</p>
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	<p>GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</p>

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4632 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4633 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4634 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-4635 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4636 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4637 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4638 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4639 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4640 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4641 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4642 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4643 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4644 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4645 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4646 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4647 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-4648 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4649 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4650 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4651 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-4652 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4653 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4654 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-4655 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-4656 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4657 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4658 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4659 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4660 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4661 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-4662 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-4663 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4664 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-4665 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-4666 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-4667 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-4668 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-4669 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-4670 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-4671 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-4672 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-4673 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-4674 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4675 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-4676 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4677 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4678 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4679 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4680 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4681 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4682 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4683 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4684 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-4685 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-4686 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

状态码： 202

表 5-4687 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.

参数	参数类型	描述
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4688 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	参数类型	描述
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4689 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4690 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.

参数	参数类型	描述
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	参数类型	描述
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

参数	参数类型	描述
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://github.com/kubernetes/enhancements/blob/master/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccountName	String	DeprecatedServiceAccountName is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	参数类型	描述
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4691 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4692 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4693 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4694 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4695 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4696 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4697 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4698 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4699 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4700 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4701 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4702 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-4703 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4704 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-4705 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4706 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4707 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4708 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4709 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4710 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4711 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4712 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-4713 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-4714 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4715 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4716 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4717 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4718 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-4719 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4720 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4721 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4722 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-4723 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4724 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4725 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4726 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4727 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4728 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4729 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4730 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-4731 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4732 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-4733 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4734 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4735 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4736 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4737 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4738 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4739 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-4740 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4741 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4742 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4743 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4744 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4745 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4746 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4747 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4748 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4749 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4750 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4751 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4752 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-4753 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4754 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4755 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4756 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-4757 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4758 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4759 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-4760 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-4761 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4762 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4763 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4764 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4765 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4766 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-4767 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-4768 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4769 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-4770 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-4771 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-4772 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-4773 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-4774 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-4775 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-4776 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-4777 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-4778 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-4779 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4780 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4781 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4782 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4783 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4784 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4785 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4786 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4787 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4788 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4789 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-4790 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-4791 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

请求示例

创建一个名为“statefulset-test”的StatefulSet负载，使用redis镜像创建3个Pod。

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "name": "statefulset-test"
  },
  "spec": {
    "replicas": 3,
    "selector": {
      "matchLabels": {
        "app": "statefulset-test"
      }
    },
    "serviceName": "statefulset-test",
    "template": {
      "metadata": {
        "labels": {
          "app": "statefulset-test"
        }
      },
      "spec": {
        "containers": [ {
          "image": "redis",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1024Mi"
            },
            "requests": {
              "cpu": "500m",
```



```
        "memory": "1024Mi"
      }
    }
  ],
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "priority": 0
}
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "creationTimestamp": "2018-09-04T07:13:00Z",
    "generation": 1,
    "labels": {
      "app": "statefulset-test"
    },
    "name": "statefulset-test",
    "namespace": "namespace-test",
    "resourceVersion": "5207623",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets/statefulset-test",
    "uid": "f4a35f35-b011-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "podManagementPolicy": "OrderedReady",
    "replicas": 3,
    "revisionHistoryLimit": 10,
    "selector": {
      "matchLabels": {
        "app": "statefulset-test"
      }
    },
    "serviceName": "",
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        }
      },
      "creationTimestamp": null,
      "labels": {
        "app": "statefulset-test"
      }
    },
    "spec": {
      "containers": [ {
        "image": "redis",
        "imagePullPolicy": "IfNotPresent",
        "name": "container-0",
        "resources": {
          "limits": {
            "cpu": "500m",
            "memory": "1Gi"
          },
          "requests": {
            "cpu": "500m",
            "memory": "1Gi"
          }
        }
      } ],
    }
  },
}
```

```
    "terminationMessagePath" : "/dev/termination-log",
    "terminationMessagePolicy" : "File"
  }],
  "dnsPolicy" : "ClusterFirst",
  "imagePullSecrets" : [ {
    "name" : "imagepull-secret"
  } ],
  "restartPolicy" : "Always",
  "schedulerName" : "default-scheduler",
  "securityContext" : { }
}
},
"updateStrategy" : {
  "type" : "OnDelete"
}
},
"status" : {
  "replicas" : 0
}
}
```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.4 删除 StatefulSet

功能介绍

删除StatefulSet。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}

表 5-4792 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the StatefulSet
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4793 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4794 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-4795 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-4796 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-4797 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4798 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4799 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-4800 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码： 202

表 5-4801 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4802 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4803 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-4804 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

- 只删除StatefulSet（对应Pod不删除）

```
{
  "Kind": "DeleteOptions",
  "apiVersion": "v1",
  "propagationPolicy": "Orphan"
}
```

- 前台级联删除（按照Pod->StatefulSet的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Foreground"
}
```

- 后台级联删除（按照StatefulSet->Pod的顺序进行删除）

```
{
  "apiVersion": "v1",
```

```
"kind": "DeleteOptions",  
"propagationPolicy": "Background"  
}
```

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "v1",  
  "code": 200,  
  "details": {  
    "group": "apps",  
    "kind": "statefulsets",  
    "name": "statefulset-test",  
    "uid": "5eb82b50-b028-11e8-9d5d-c88d83be759f"  
  },  
  "kind": "Status",  
  "metadata": { },  
  "status": "Success"  
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.5 查询 StatefulSet

功能介绍

查询StatefulSet的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}

表 5-4805 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the StatefulSet
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4806 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4807 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-4808 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4809 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.

参数	参数类型	描述
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4810 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4811 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4812 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4813 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4814 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4815 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4816 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4817 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4818 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4819 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4820 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4821 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4822 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4823 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-4824 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4825 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-4826 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4827 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4828 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4829 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4830 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4831 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4832 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4833 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-4834 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-4835 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4836 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4837 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4838 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4839 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-4840 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4841 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4842 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4843 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-4844 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4845 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4846 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4847 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4848 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4849 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4850 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4851 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-4852 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4853 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-4854 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4855 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4856 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4857 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4858 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4859 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4860 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-4861 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4862 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4863 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4864 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4865 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4866 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4867 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4868 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4869 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4870 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4871 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4872 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4873 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-4874 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4875 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4876 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4877 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-4878 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4879 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4880 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-4881 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-4882 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4883 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4884 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4885 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4886 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4887 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-4888 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-4889 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4890 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-4891 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-4892 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-4893 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-4894 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-4895 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-4896 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-4897 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-4898 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-4899 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-4900 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-4901 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4902 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-4903 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-4904 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-4905 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-4906 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-4907 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-4908 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-4909 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-4910 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-4911 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-4912 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "creationTimestamp": "2018-09-04T07:13:00Z",
    "generation": 1,
    "labels": {
      "app": "statefulset-test"
    },
    "name": "statefulset-test",
    "namespace": "namespace-test",
    "resourceVersion": "5209881",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets/statefulset-test",
    "uid": "f4a35f35-b011-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "podManagementPolicy": "OrderedReady",
    "replicas": 3,
    "revisionHistoryLimit": 10,
    "selector": {
      "matchLabels": {
        "app": "statefulset-test"
      }
    }
  }
}
```

```

    },
    "serviceName": "",
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "app": "statefulset-test"
        }
      },
      "spec": {
        "containers": [ {
          "image": "redis",
          "imagePullPolicy": "IfNotPresent",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1Gi"
            },
            "requests": {
              "cpu": "500m",
              "memory": "1Gi"
            }
          },
          "terminationMessagePath": "/dev/termination-log",
          "terminationMessagePolicy": "File"
        } ],
        "dnsPolicy": "ClusterFirst",
        "imagePullSecrets": [ {
          "name": "imagepull-secret"
        } ],
        "restartPolicy": "Always",
        "schedulerName": "default-scheduler",
        "securityContext": { }
      }
    },
    "updateStrategy": {
      "type": "OnDelete"
    },
    "status": {
      "collisionCount": 0,
      "currentReplicas": 3,
      "currentRevision": "statefulset-test-f986b645b",
      "observedGeneration": 1,
      "readyReplicas": 2,
      "replicas": 3,
      "updateRevision": "statefulset-test-f986b645b"
    }
  }
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.6 更新 StatefulSet

功能介绍

更新StatefulSet。

The following fields can be updated:

- metadata.labels
- metadata.annotations
- spec.replicas
- spec.template
- spec.restartPolicy
- spec.revisionHistoryLimit
- spec.progressDeadlineSeconds The other fields cannot be updated.

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}

表 5-4913 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the StatefulSet

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-4914 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-4915 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	目前支持三种类型的PATCH请求方法的操作，参考《 使用JSON合并patch更新Deployment 》。 <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-4916 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-4917 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-4918 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.

参数	参数类型	描述
updateStrategy	io.k8s.api.app.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-4919 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-4920 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-4921 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-4922 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-4923 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-4924 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-4925 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-4926 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-4927 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4928 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-4929 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-4930 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-4931 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-4932 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-4933 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4934 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-4935 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-4936 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-4937 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-4938 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-4939 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-4940 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4941 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-4942 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-4943 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-4944 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-4945 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-4946 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-4947 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4948 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-4949 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-4950 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-4951 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-4952 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-4953 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-4954 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-4955 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-4956 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-4957 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-4958 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-4959 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4960 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-4961 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-4962 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-4963 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-4964 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-4965 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-4966 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-4967 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-4968 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-4969 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-4970 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-4971 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-4972 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4973 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-4974 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-4975 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-4976 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-4977 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-4978 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-4979 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-4980 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-4981 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-4982 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-4983 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-4984 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-4985 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-4986 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-4987 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-4988 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-4989 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-4990 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-4991 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-4992 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-4993 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-4994 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-4995 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-4996 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-4997 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-4998 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-4999 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5000 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5001 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5002 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5003 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5004 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5005 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5006 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5007 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-5008 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-5009 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-5010 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-5011 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5012 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5013 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5014 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5015 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5016 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5017 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5018 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-5019 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-5020 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-5021 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

请求示例

更新StatefulSet中的labels值为"app": "statefulset-test2"。

```
{
  "metadata": {
    "labels": {
      "app": "statefulset-test2"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "creationTimestamp": "2018-09-04T07:13:00Z",
    "generation": 2,
    "labels": {
      "app": "statefulset-test2"
    }
  },
  "name": "statefulset-test",
  "namespace": "namespace-test",
  "resourceVersion": "5231369",
  "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets/statefulset-test",
  "uid": "f4a35f35-b011-11e8-b6ef-f898ef6c78b4"
}
```

```
"spec" : {
  "podManagementPolicy" : "OrderedReady",
  "replicas" : 2,
  "revisionHistoryLimit" : 10,
  "selector" : {
    "matchLabels" : {
      "app" : "statefulset-test"
    }
  },
  "serviceName" : "",
  "template" : {
    "metadata" : {
      "creationTimestamp" : null,
      "labels" : {
        "app" : "statefulset-test"
      }
    },
    "spec" : {
      "containers" : [ {
        "image" : "redis",
        "imagePullPolicy" : "IfNotPresent",
        "name" : "container-0",
        "resources" : {
          "limits" : {
            "cpu" : "500m",
            "memory" : "1Gi"
          },
          "requests" : {
            "cpu" : "500m",
            "memory" : "1Gi"
          }
        },
        "terminationMessagePath" : "/dev/termination-log",
        "terminationMessagePolicy" : "File"
      } ],
      "dnsPolicy" : "ClusterFirst",
      "imagePullSecrets" : [ {
        "name" : "imagepull-secret"
      } ],
      "restartPolicy" : "Always",
      "schedulerName" : "default-scheduler",
      "securityContext" : { }
    }
  },
  "updateStrategy" : {
    "type" : "OnDelete"
  }
},
"status" : {
  "collisionCount" : 0,
  "currentRevision" : "statefulset-test-f986b645b",
  "observedGeneration" : 2,
  "replicas" : 2,
  "updateRevision" : "statefulset-test-7748d5459",
  "updatedReplicas" : 2
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized

状态码	描述
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.7 替换 StatefulSet

功能介绍

替换StatefulSet。

The following fields can be updated:

- metadata.labels
- metadata.annotations
- spec.template
- spec.replicas
- spec.revisionHistoryLimit
- spec.progressDeadlineSeconds The other fields cannot be updated.

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}

表 5-5022 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the StatefulSet

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-5023 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-5024 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-5025 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	否	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	否	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-5026 io.k8s.api.apps.v1.StatefulSetSpec

参数	是否必选	参数类型	描述
podManagementPolicy	否	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	否	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	否	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	是	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	是否必选	参数类型	描述
serviceName	是	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	是	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	否	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	否	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-5027 io.k8s.api.core.v1.PodTemplateSpec

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5028 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.

参数	是否必选	参数类型	描述
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	是否必选	参数类型	描述
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod

参数	是否必选	参数类型	描述
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	是否必选	参数类型	描述
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.

参数	是否必选	参数类型	描述
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy

参数	是否必选	参数类型	描述
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	否	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	是否必选	参数类型	描述
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	否	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.

参数	是否必选	参数类型	描述
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5029 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	是否必选	参数类型	描述
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5030 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5031 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5032 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5033 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5034 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5035 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-5036 io.k8s.api.core.v1.PodAntiAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-5037 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5038 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5039 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5040 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.
value	否	String	value is the value of the option

表 5-5041 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	是否必选	参数类型	描述
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5042 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-5043 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5044 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5045 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5046 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5047 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5048 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-5049 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5050 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-5051 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-5052 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5053 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5054 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5055 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5056 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-5057 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5058 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5059 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5060 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-5061 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5062 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5063 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain ':'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5064 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5065 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5066 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5067 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5068 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-5069 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5070 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-5071 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
gmsaCredentialSpecName	否	String	<code>GMSACredentialSpecName</code> is the name of the GMSA credential spec to use.
runAsUserName	否	String	The <code>UserName</code> in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in <code>PodSecurityContext</code> . If set in both <code>SecurityContext</code> and <code>PodSecurityContext</code> , the value specified in <code>SecurityContext</code> takes precedence.

表 5-5072 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are <code>NoSchedule</code> , <code>PreferNoSchedule</code> and <code>NoExecute</code> .
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be <code>Exists</code> ; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5073 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5074 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5075 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5076 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5077 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-5078 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5079 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5080 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5081 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5082 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5083 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5084 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5085 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5086 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5087 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5088 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5089 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5090 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-5091 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5092 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5093 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.

参数	是否必选	参数类型	描述
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5094 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or</p>

参数	是否必选	参数类型	描述
			<p>suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-5095 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5096 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5097 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-5098 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-5099 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5100 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5101 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5102 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5103 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5104 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-5105 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-5106 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-5107 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5108 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5109 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5110 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5111 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5112 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5113 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5114 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmdk

表 5-5115 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	是否必选	参数类型	描述
rollingUpdate	否	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	否	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-5116 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	是否必选	参数类型	描述
partition	否	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-5117 io.k8s.api.core.v1.PersistentVolumeClaim

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	否	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-5118 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.
labels	否	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5119 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5120 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5121 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5122 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-5123 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5124 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5125 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5126 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	否	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	否	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	否	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-5127 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time we probed the condition.
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
message	否	String	Human-readable message indicating details about last transition.
reason	否	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	是	String	status is the status of the condition.
type	是	String	type is the type of the condition.

表 5-5128 io.k8s.api.apps.v1.StatefulSetStatus

参数	是否必选	参数类型	描述
collisionCount	否	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	否	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	否	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	否	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).

参数	是否必选	参数类型	描述
observedGeneration	否	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	否	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	是	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	否	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)
updatedReplicas	否	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-5129 io.k8s.api.apps.v1.StatefulSetCondition

参数	是否必选	参数类型	描述
lastTransitionTime	否	String	Last time the condition transitioned from one status to another.
message	否	String	A human readable message indicating details about the transition.
reason	否	String	The reason for the condition's last transition.
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of statefulset condition.

响应参数

状态码： 200

表 5-5130 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-5131 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.

参数	参数类型	描述
updateStrategy	io.k8s.api.app.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-5132 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5133 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5134 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5135 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5136 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5137 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5138 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5139 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5140 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5141 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5142 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5143 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5144 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5145 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5146 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5147 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5148 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5149 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5150 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5151 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5152 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5153 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5154 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5155 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5156 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5157 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5158 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5159 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5160 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5161 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5162 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5163 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5164 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5165 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5166 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5167 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5168 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5169 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5170 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5171 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5172 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5173 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5174 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5175 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5176 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5177 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5178 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5179 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	<p>FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.</p>
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	<p>Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running</p>
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	<p>GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</p>

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5180 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5181 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5182 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5183 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5184 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5185 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5186 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5187 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5188 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5189 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5190 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5191 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5192 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5193 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5194 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5195 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5196 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5197 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5198 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5199 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-5200 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5201 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5202 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5203 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5204 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5205 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5206 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5207 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5208 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5209 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5210 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5211 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5212 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5213 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5214 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5215 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5216 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5217 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5218 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5219 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5220 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-5221 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-5222 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-5223 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5224 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5225 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5226 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5227 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5228 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/

表 5-5229 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5230 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5231 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-5232 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-5233 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-5234 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

状态码： 201

表 5-5235 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.

参数	参数类型	描述
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-5236 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	参数类型	描述
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-5237 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5238 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.

参数	参数类型	描述
hostAliases	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	参数类型	描述
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.

参数	参数类型	描述
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://github.com/kubernetes/enhancements/blob/master/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	参数类型	描述
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5239 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5240 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5241 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5242 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5243 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5244 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5245 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5246 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5247 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5248 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5249 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5250 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5251 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5252 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5253 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5254 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5255 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5256 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5257 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5258 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5259 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5260 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5261 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5262 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5263 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5264 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5265 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5266 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5267 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5268 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5269 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5270 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5271 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5272 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5273 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain '%'. Must not contain '*'. Must not contain '?'. Must not contain '['. Must not contain ']'. Must not contain '&'. Must not contain '='. Must not contain '". Must not contain ''.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5274 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5275 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5276 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5277 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5278 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5279 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5280 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5281 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5282 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5283 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5284 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5285 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5286 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5287 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5288 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5289 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5290 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5291 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5292 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5293 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5294 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5295 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5296 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5297 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5298 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5299 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5300 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5301 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5302 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5303 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5304 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-5305 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5306 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5307 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5308 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5309 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5310 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5311 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5312 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5313 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5314 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5315 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5316 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5317 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5318 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5319 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5320 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5321 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5322 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5323 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5324 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5325 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-5326 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-5327 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-5328 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5329 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5330 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5331 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5332 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5333 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5334 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5335 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5336 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-5337 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-5338 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-5339 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

请求示例

将StatefulSet中的replicas值修改为2。

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "name": "statefulset-test"
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "matchLabels": {
        "app": "statefulset-test"
      }
    },
    "serviceName": "statefulset-test",
    "template": {
      "metadata": {
        "labels": {
          "app": "statefulset-test"
        }
      },
      "spec": {
        "containers": [ {
          "image": "k8s.gcr.io/redis:20202/ccr/redis:v1",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1024Mi"
            },
            "requests": {
              "cpu": "500m",
```

```
        "memory": "1024Mi"
      }
    }
  ],
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ]
}
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "creationTimestamp": "2018-09-04T07:13:00Z",
    "generation": 2,
    "labels": {
      "app": "statefulset-test"
    },
    "name": "statefulset-test",
    "namespace": "namespace-test",
    "resourceVersion": "5223616",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets/statefulset-test",
    "uid": "f4a35f35-b011-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "podManagementPolicy": "OrderedReady",
    "replicas": 2,
    "revisionHistoryLimit": 10,
    "selector": {
      "matchLabels": {
        "app": "statefulset-test"
      }
    },
    "serviceName": "",
    "template": {
      "metadata": {
        "creationTimestamp": null,
        "labels": {
          "app": "statefulset-test"
        }
      },
      "spec": {
        "containers": [ {
          "image": "redis",
          "imagePullPolicy": "IfNotPresent",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1Gi"
            },
            "requests": {
              "cpu": "500m",
              "memory": "1Gi"
            }
          }
        },
        "terminationMessagePath": "/dev/termination-log",
        "terminationMessagePolicy": "File"
      }
    },
    "dnsPolicy": "ClusterFirst",
```



```
"imagePullSecrets" : [ {  
  "name" : "imagepull-secret"  
} ],  
"restartPolicy" : "Always",  
"schedulerName" : "default-scheduler",  
"securityContext" : { }  
}  
},  
"updateStrategy" : {  
  "type" : "OnDelete"  
}  
},  
"status" : {  
  "collisionCount" : 0,  
  "currentReplicas" : 3,  
  "currentRevision" : "statefulset-test-f986b645b",  
  "observedGeneration" : 1,  
  "readyReplicas" : 1,  
  "replicas" : 3,  
  "updateRevision" : "statefulset-test-f986b645b"  
}  
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.8 查询 StatefulSet 状态

功能介绍

查询StatefulSet状态。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status

表 5-5340 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the StatefulSet
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-5341 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-5342 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-5343 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-5344 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.

参数	参数类型	描述
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-5345 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5346 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5347 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5348 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5349 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5350 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5351 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5352 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5353 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5354 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5355 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5356 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5357 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5358 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5359 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5360 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5361 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5362 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5363 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5364 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5365 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5366 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5367 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5368 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5369 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5370 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5371 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5372 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5373 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5374 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5375 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5376 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5377 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5378 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5379 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5380 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5381 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5382 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5383 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5384 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5385 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5386 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5387 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5388 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5389 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5390 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5391 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5392 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5393 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5394 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5395 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5396 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5397 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5398 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5399 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5400 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5401 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5402 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5403 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5404 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5405 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5406 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5407 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5408 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5409 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5410 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5411 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5412 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-5413 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5414 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5415 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5416 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5417 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5418 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5419 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5420 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5421 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5422 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5423 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5424 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5425 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5426 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5427 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5428 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5429 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5430 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5431 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5432 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5433 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-5434 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-5435 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-5436 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5437 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5438 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5439 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5440 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5441 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5442 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5443 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5444 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-5445 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-5446 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-5447 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "StatefulSet",
  "metadata": {
    "creationTimestamp": "2018-09-04T07:13:00Z",
    "generation": 1,
    "labels": {
      "app": "statefulset-test"
    },
    "name": "statefulset-test",
    "namespace": "namespace-test",
    "resourceVersion": "5217947",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/statefulsets/statefulset-test/status",
    "uid": "f4a35f35-b011-11e8-b6ef-f898ef6c78b4"
  },
  "spec": {
    "podManagementPolicy": "OrderedReady",
    "replicas": 3,
    "revisionHistoryLimit": 10,
    "selector": {
      "matchLabels": {
        "app": "statefulset-test"
      }
    }
  }
}
```

```

    },
    "serviceName": "",
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "app": "statefulset-test"
        }
      },
      "spec": {
        "containers": [ {
          "image": "*/*/*/20202/cci/redis:V1",
          "imagePullPolicy": "IfNotPresent",
          "name": "container-0",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1Gi"
            },
            "requests": {
              "cpu": "500m",
              "memory": "1Gi"
            }
          },
          "terminationMessagePath": "/dev/termination-log",
          "terminationMessagePolicy": "File"
        } ],
        "dnsPolicy": "ClusterFirst",
        "imagePullSecrets": [ {
          "name": "imagepull-secret"
        } ],
        "restartPolicy": "Always",
        "schedulerName": "default-scheduler",
        "securityContext": { }
      }
    },
    "updateStrategy": {
      "type": "OnDelete"
    },
    "status": {
      "collisionCount": 0,
      "currentReplicas": 3,
      "currentRevision": "statefulset-test-f986b645b",
      "observedGeneration": 1,
      "readyReplicas": 1,
      "replicas": 3,
      "updateRevision": "statefulset-test-f986b645b"
    }
  }
}

```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound

状态码	描述
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.19.9 查询用户所有的 StatefulSets

功能介绍

This API is used to list all StatefulSet resource objects.

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/statefulsets

表 5-5448 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 5-5449 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-5450 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.apps.v1.StatefulSet objects	StatefulSet represents a set of pods with consistent identities. Identities are defined as: <ul style="list-style-type: none">• Network: A single stable DNS and hostname.• Storage: As many VolumeClaims as requested. The StatefulSet guarantees that a given network identity will always map to the same storage identity.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	ListMeta describes metadata that synthetic resources must have, including lists and various status objects. A resource may have only one of {ObjectMeta, ListMeta}.

表 5-5451 io.k8s.api.apps.v1.StatefulSet

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	ObjectMeta is metadata that all persisted resources must have, which includes all objects users must create.
spec	io.k8s.api.apps.v1.StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	io.k8s.api.apps.v1.StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 5-5452 io.k8s.api.apps.v1.StatefulSetSpec

参数	参数类型	描述
podManagementPolicy	String	podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down. The default policy is <i>OrderedReady</i> , where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order. The alternative policy is <i>Parallel</i> which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.
replicas	Integer	replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1.

参数	参数类型	描述
revisionHistoryLimit	Integer	revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	selector is a label query over pods that should match the replica count. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
serviceName	String	serviceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.
template	io.k8s.api.core.v1.PodTemplateSpec object	template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.
updateStrategy	io.k8s.api.apps.v1.StatefulSetUpdateStrategy object	updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.
volumeClaimTemplates	Array of io.k8s.api.core.v1.PersistentVolumeClaim objects	volumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name.

表 5-5453 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5454 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5455 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5456 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5457 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5458 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5459 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5460 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5461 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5462 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5463 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5464 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5465 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5466 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5467 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5468 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5469 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5470 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5471 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5472 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5473 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5474 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5475 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5476 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5477 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5478 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5479 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5480 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5481 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5482 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5483 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5484 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5485 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5486 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5487 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5488 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5489 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain ':'. Must not contain '*'. Must not contain '?'. Must not contain '['. Must not contain ']'. Must not contain ' '. Must not contain '&'. Must not contain '='. Must not contain '". Must not contain ''.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5490 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5491 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5492 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5493 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5494 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5495 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5496 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5497 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5498 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5499 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5500 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5501 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5502 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5503 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5504 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5505 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5506 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5507 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5508 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5509 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5510 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5511 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5512 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5513 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5514 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5515 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5516 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5517 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5518 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5519 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5520 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-5521 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5522 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5523 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5524 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5525 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5526 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5527 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5528 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5529 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5530 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5531 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5532 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5533 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5534 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5535 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5536 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5537 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5538 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5539 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5540 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5541 io.k8s.api.apps.v1.StatefulSetUpdateStrategy

参数	参数类型	描述
rollingUpdate	io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.
type	String	Type indicates the type of the StatefulSetUpdateStrategy. Default is RollingUpdate.

表 5-5542 io.k8s.api.apps.v1.RollingUpdateStatefulSetStrategy

参数	参数类型	描述
partition	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned. Default value is 0.

表 5-5543 io.k8s.api.core.v1.PersistentVolumeClaim

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
status	io.k8s.api.core.v1.PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims

表 5-5544 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations

参数	参数类型	描述
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5545 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5546 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5547 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5548 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5549 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5550 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5551 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5552 io.k8s.api.core.v1.PersistentVolumeClaimStatus

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1
capacity	Map<String,String>	Represents the actual resources of the underlying volume.
conditions	Array of io.k8s.api.core.v1.PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to 'ResizeStarted'.
phase	String	Phase represents the current phase of PersistentVolumeClaim.

表 5-5553 io.k8s.api.core.v1.PersistentVolumeClaimCondition

参数	参数类型	描述
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	Human-readable message indicating details about last transition.

参数	参数类型	描述
reason	String	Unique, this should be a short, machine understandable string that gives the reason for condition's last transition. If it reports "ResizeStarted" that means the underlying persistent volume is being resized.
status	String	status is the status of the condition.
type	String	type is the type of the condition.

表 5-5554 io.k8s.api.apps.v1.StatefulSetStatus

参数	参数类型	描述
collisionCount	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision.
conditions	Array of io.k8s.api.apps.v1.StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.
currentReplicas	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).
observedGeneration	Long	observedGeneration is the most recent generation observed for this StatefulSet. It corresponds to the StatefulSet's generation, which is updated on mutation by the API Server.
readyReplicas	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
replicas	Integer	replicas is the number of Pods created by the StatefulSet controller.
updateRevision	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)

参数	参数类型	描述
updatedReplicas	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.

表 5-5555 io.k8s.api.apps.v1.StatefulSetCondition

参数	参数类型	描述
lastTransitionTime	String	Last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of statefulset condition.

表 5-5556 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1beta1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2017-04-18T06:05:02Z",
      "generateName": "sz",
      "generation": 1,
      "labels": {
```

```
    "app" : "mysql"
  },
  "name" : "mysql",
  "namespace" : "default",
  "resourceVersion" : "1809843",
  "selfLink" : "/apis/apps/v1beta1/namespaces/default/statefulsets/mysql",
  "uid" : "f5cf50f5-23fc-11e7-9c83-fa163ec08232"
},
"spec" : {
  "replicas" : 1,
  "selector" : {
    "matchLabels" : {
      "app" : "mysql"
    }
  }
},
"serviceName" : "mysql-service",
"template" : {
  "metadata" : {
    "creationTimestamp" : null,
    "labels" : {
      "app" : "mysql"
    },
    "name" : "-sz"
  },
  "spec" : {
    "containers" : [ {
      "image" : "10.154.52.159:443/test/nginx:latest",
      "imagePullPolicy" : "IfNotPresent",
      "name" : "container01",
      "ports" : [ {
        "containerPort" : 80,
        "protocol" : "TCP"
      } ],
      "resources" : { },
      "terminationMessagePath" : "/dev/termination-log"
    } ],
    "dnsPolicy" : "ClusterFirst",
    "restartPolicy" : "Always",
    "securityContext" : { }
  }
},
"volumeClaimTemplates" : [ {
  "metadata" : {
    "creationTimestamp" : null,
    "name" : "db"
  },
  "spec" : {
    "accessModes" : [ "ReadWriteOnce" ],
    "resources" : {
      "requests" : {
        "storage" : "1Gi"
      }
    }
  }
},
  "status" : {
    "phase" : "Pending"
  }
} ]
},
"status" : {
  "replicas" : 1
}
} ],
"kind" : "StatefulSetList",
"metadata" : {
  "resourceVersion" : "1809953",
  "selfLink" : "/apis/apps/v1beta1/statefulsets"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20 Job

5.20.1 查询用户所有的 Jobs

功能介绍

This API is used to obtain a Job list.

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1/jobs

表 5-5557 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
pretty	否	String	If 'true', then the output is pretty printed.

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求参数

表 5-5558 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-5559 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.batch.v1.Job objects	items is the list of Jobs.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-5560 io.k8s.api.batch.v1.Job

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5561 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((.spec.completions - .status.successful) < .spec.parallelism)$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-5562 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5563 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5564 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5565 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5566 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5567 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5568 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5569 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5570 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5571 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5572 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5573 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5574 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5575 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5576 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5577 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5578 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5579 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5580 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5581 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5582 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5583 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5584 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5585 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5586 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5587 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5588 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5589 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5590 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5591 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5592 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5593 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5594 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5595 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5596 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5597 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5598 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain ':'. Must not contain '*'. Must not contain '?'. Must not contain '['. Must not contain ']'. Must not contain '&'. Must not contain '='. Must not contain '". Must not contain ''.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5599 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5600 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5601 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5602 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5603 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5604 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5605 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5606 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entryptpoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5607 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5608 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5609 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5610 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5611 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5612 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5613 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5614 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5615 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5616 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5617 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5618 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5619 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5620 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5621 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5622 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5623 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5624 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5625 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5626 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5627 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5628 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5629 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5630 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5631 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5632 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5633 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5634 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5635 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5636 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5637 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-5638 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5639 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5640 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5641 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5642 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5643 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5644 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5645 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5646 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5647 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5648 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5649 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5650 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5651 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5652 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5653 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5654 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5655 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5656 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5657 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5658 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-5659 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

表 5-5660 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2017-12-13T03:15:55Z",
      "labels": {
        "name": "job-test"
      },
      "name": "jobs-12130306",
      "namespace": "ns-12130306-s",
      "resourceVersion": "419064",
      "selfLink": "/apis/batch/v1/namespaces/ns-12130306-s/jobs/jobs-12130306",
      "uid": "eed6b02b-dfb3-11e7-9c19-fa163e2d897b"
    },
    "spec": {
      "completions": 1,
      "parallelism": 1,
      "selector": {
        "matchLabels": {
          "controller-uid": "eed6b02b-dfb3-11e7-9c19-fa163e2d897b"
        }
      },
      "template": {
        "metadata": {
          "creationTimestamp": null,
          "labels": {
            "controller-uid": "eed6b02b-dfb3-11e7-9c19-fa163e2d897b",
            "job-name": "jobs-12130306",
            "name": "job-test"
          }
        },

```



```
    "name" : "jobs-12130306"
  },
  "spec" : {
    "containers" : [ {
      "image" : "172.16.5.235:20202/test/redis:latest",
      "imagePullPolicy" : "Always",
      "name" : "jobs-12130306",
      "resources" : { },
      "terminationMessagePath" : "/dev/termination-log",
      "terminationMessagePolicy" : "File"
    } ],
    "dnsPolicy" : "ClusterFirst",
    "restartPolicy" : "Never",
    "schedulerName" : "default-scheduler",
    "securityContext" : { }
  }
},
"status" : {
  "active" : 1,
  "startTime" : "2017-12-13T03:15:55Z"
}
}],
"kind" : "JobList",
"metadata" : {
  "resourceVersion" : "419065",
  "selfLink" : "/apis/batch/v1/jobs"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.2 删除指定 namespace 下的 Jobs

功能介绍

删除Namespace下所有Job。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/batch/v1/namespaces/{namespace}/jobs

表 5-5661 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-5662 Query 参数

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset

参数	是否必选	参数类型	描述
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-5663 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-5664 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-5665 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-5666 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5667 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5668 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-5669 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

- 只删除Job（对应Pod不删除）。

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions",
  "propagationPolicy": "Orphan"
}
```

- 前台级联删除（按照Pod->Job的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Foreground"
}
```

- 后台级联删除（按照Job->Pod的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Background"
}
```

响应示例

状态码： 200

OK

```
{
  "apiVersion": "batch/v1",
  "items": null,
  "kind": "JobList",
  "metadata": {
    "resourceVersion": "5415316",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs"
  }
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.3 查询指定 namespace 下的 Jobs

功能介绍

查询Namespace下所有Job的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1/namespaces/{namespace}/jobs

表 5-5670 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-5671 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-5672 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-5673 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.batch.v1.Job objects	items is the list of Jobs.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata

表 5-5674 io.k8s.api.batch.v1.Job

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5675 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((.spec.completions - .status.successful) < .spec.parallelism)$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-5676 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5677 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5678 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5679 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5680 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5681 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5682 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5683 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5684 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5685 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5686 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5687 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5688 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5689 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5690 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5691 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5692 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5693 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5694 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5695 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5696 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5697 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5698 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5699 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5700 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5701 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5702 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5703 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, 0 < x < 65536.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5704 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5705 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5706 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5707 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5708 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5709 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5710 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5711 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5712 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'. Must not contain '..'. Must not contain '.'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5713 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5714 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5715 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5716 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5717 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5718 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5719 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5720 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5721 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5722 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5723 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5724 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5725 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5726 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5727 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5728 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5729 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5730 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5731 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5732 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5733 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5734 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5735 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-5736 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5737 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5738 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5739 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5740 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5741 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5742 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5743 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5744 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5745 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5746 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5747 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5748 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5749 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5750 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5751 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-5752 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5753 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5754 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5755 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5756 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5757 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5758 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5759 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5760 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5761 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5762 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5763 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5764 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5765 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5766 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5767 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5768 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5769 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5770 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5771 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5772 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-5773 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

表 5-5774 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "batch/v1",
  "items": [ {
    "metadata": {
      "creationTimestamp": "2018-09-05T01:10:59Z",
      "labels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
        "job-name": "pi"
      },
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5391205",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",
    "uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
      }
    },
  },
  "template": {
    "metadata": {
      "annotations": {
        "cri.cci.io/container-type": "secure-container"
      },
    },
    "creationTimestamp": null,
```

```
"labels": {
  "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
  "job-name": "pi"
},
"name": "pi"
},
"spec": {
  "containers": [ {
    "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
    "image": "perl",
    "imagePullPolicy": "Always",
    "name": "pi",
    "resources": {
      "limits": {
        "cpu": "500m",
        "memory": "1Gi"
      },
      "requests": {
        "cpu": "500m",
        "memory": "1Gi"
      }
    }
  },
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
} ],
"dnsPolicy": "ClusterFirst",
"imagePullSecrets": [ {
  "name": "imagepull-secret"
} ],
"restartPolicy": "Never",
"schedulerName": "default-scheduler",
"securityContext": { }
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}],
"kind": "JobList",
"metadata": {
  "resourceVersion": "5391600",
  "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict

状态码	描述
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.4 创建 Job

功能介绍

创建Job。

说明

Note: Kubernetes 只有在 Pod 结束 (Terminated) 的时候才会发送 preStop 事件, 这意味着在 Pod 完成 (Completed) 时 preStop 的事件处理逻辑不会被触发。这个限制在 [issue #55087](#) 中被追踪。

调用方法

请参见[如何调用API](#)。

URI

POST /apis/batch/v1/namespaces/{namespace}/jobs

表 5-5775 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-5776 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-5777 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-5778 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5779 io.k8s.api.batch.v1.JobSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	否	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	否	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	是否必选	参数类型	描述
manualSelector	否	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	否	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	是否必选	参数类型	描述
template	是	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/job-run-to-completion/
tTlSecondsAfterFinished	否	Integer	tTlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, tTlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-5780 io.k8s.api.core.v1.PodTemplateSpec

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5781 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.

参数	是否必选	参数类型	描述
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.

参数	是否必选	参数类型	描述
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	是否必选	参数类型	描述
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.

参数	是否必选	参数类型	描述
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	是否必选	参数类型	描述
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceName	否	String	DeprecatedServiceAccount is a depreciated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.

参数	是否必选	参数类型	描述
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.

参数	是否必选	参数类型	描述
subdomain	否	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	是否必选	参数类型	描述
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5782 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5783 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5784 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	是否必选	参数类型	描述
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5785 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5786 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5787 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	是否必选	参数类型	描述
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5788 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-5789 `io.k8s.api.core.v1.PodAntiAffinity`

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5790 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5791 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"

参数	是否必选	参数类型	描述
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5792 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5793 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.

参数	是否必选	参数类型	描述
value	否	String	value is the value of the option

表 5-5794 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	是否必选	参数类型	描述
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	是否必选	参数类型	描述
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5795 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-5796 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5797 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5798 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5799 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5800 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5801 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-5802 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5803 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-5804 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-5805 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5806 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5807 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5808 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5809 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-5810 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5811 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5812 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5813 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-5814 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5815 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5816 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain ':'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5817 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5818 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5819 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5820 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5821 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-5822 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5823 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-5824 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	否	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	否	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5825 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5826 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5827 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5828 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5829 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5830 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-5831 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5832 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5833 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5834 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5835 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5836 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5837 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5838 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5839 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5840 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5841 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5842 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5843 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-5844 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5845 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5846 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5847 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5848 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5849 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5850 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5851 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-5852 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5853 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5854 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.

参数	是否必选	参数类型	描述
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5855 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) (Note that [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) may be empty, from the "" case in [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br).) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= 0 1 ... 9 [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br). ./[topic/body/section/</p>

参数	是否必选	参数类型	描述
			<p>table/tgroup/tbody/row/entry/p/br {""} (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "+" "-" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2^63-1</p>

参数	是否必选	参数类型	描述
			<p>in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-5856 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5857 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5858 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-5859 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-5860 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5861 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5862 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5863 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5864 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5865 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-5866 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-5867 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-5868 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5869 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5870 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5871 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5872 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5873 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5874 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5875 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmk

表 5-5876 io.k8s.api.batch.v1.JobStatus

参数	是否必选	参数类型	描述
active	否	Integer	The number of actively running pods.

参数	是否必选	参数类型	描述
completionTime	否	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	否	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	否	Integer	The number of pods which reached phase Failed.
startTime	否	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	否	Integer	The number of pods which reached phase Succeeded.

表 5-5877 io.k8s.api.batch.v1.JobCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time the condition was checked.
lastTransitionTime	否	String	Last time the condition transit from one status to another.
message	否	String	Human readable message indicating details about last transition.
reason	否	String	(brief) reason for the condition's last transition.
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of job condition, Complete or Failed.

响应参数

状态码： 200

表 5-5878 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5879 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer

参数	参数类型	描述
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((.spec.completions - .status.successful) < .spec.parallelism)$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
<code>ttlSecondsAfterFinished</code>	Integer	<code>ttlSecondsAfterFinished</code> limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, <code>ttlSecondsAfterFinished</code> after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the <code>TTLAfterFinished</code> feature.

表 5-5880 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
<code>metadata</code>	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
<code>spec</code>	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5881 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
<code>activeDeadlineSeconds</code>	Long	Optional duration in seconds the pod may be active on the node relative to <code>StartTime</code> before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
<code>affinity</code>	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
<code>automountServiceAccountToken</code>	Boolean	<code>AutomountServiceAccountToken</code> indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5882 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5883 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5884 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5885 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5886 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5887 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5888 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5889 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5890 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5891 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5892 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5893 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5894 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5895 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5896 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5897 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5898 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5899 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-5900 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-5901 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5902 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-5903 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-5904 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-5905 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-5906 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-5907 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-5908 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5909 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-5910 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-5911 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-5912 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-5913 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-5914 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-5915 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-5916 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '..'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-5917 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-5918 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-5919 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-5920 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5921 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-5922 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-5923 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-5924 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-5925 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-5926 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-5927 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-5928 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-5929 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-5930 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-5931 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-5932 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-5933 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5934 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-5935 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-5936 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-5937 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-5938 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-5939 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-5940 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-5941 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-5942 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-5943 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-5944 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-5945 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-5946 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-5947 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-5948 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-5949 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-5950 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-5951 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-5952 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-5953 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-5954 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-5955 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) (Note that [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) may be empty, from the "" case in [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= 0 1 ... 9 [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br). .[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= "+" "-" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br)[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}] (br) ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) [/topic/body/section/table/tgroup/</p>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-5956 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-5957 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-5958 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-5959 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-5960 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-5961 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-5962 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-5963 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-5964 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-5965 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-5966 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-5967 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-5968 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-5969 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-5970 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-5971 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-5972 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-5973 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-5974 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-5975 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-5976 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-5977 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

状态码： 201

表 5-5978 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5979 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((\text{.spec.completions} - \text{.status.successful}) < \text{.spec.parallelism})$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	参数类型	描述
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-5980 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-5981 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br).[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br).svc.[/topic/body/section/table/tgroup/tbody/row/entry/p/br {}"] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-5982 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-5983 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-5984 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-5985 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-5986 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-5987 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-5988 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5989 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-5990 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-5991 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-5992 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-5993 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-5994 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5995 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-5996 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-5997 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-5998 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-5999 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6000 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6001 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6002 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6003 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6004 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6005 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6006 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6007 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6008 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6009 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6010 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6011 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6012 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-6013 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-6014 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-6015 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-6016 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '/'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6017 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6018 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6019 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6020 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6021 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6022 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6023 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6024 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6025 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6026 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6027 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6028 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6029 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6030 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6031 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6032 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6033 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6034 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6035 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6036 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6037 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6038 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6039 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6040 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6041 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6042 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6043 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6044 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6045 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6046 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6047 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6048 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6049 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6050 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6051 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6052 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6053 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6054 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6055 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre> [/<topic ""="" "-"="" (br)="" (br).="" (br).)="" (br)[="" (international="" (note="" ...="" 1="" 9="" ::="Ki" <a="" <topic="" [="" be="" body="" br="" case="" ei="" empty,="" entry="" from="" gi="" href="http://physics.nist.gov/cuu/Units/binary.html" in="" may="" mi="" of="" p="" pi="" row="" section="" see:="" system="" table="" tbody="" tgroup="" that="" the="" ti="" units;="" {""}="" ="">http://physics.nist.gov/cuu/Units/binary.html) [/<topic <="" body="" pre="" section="" table="" tgroup=""> </topic></topic></pre>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-6056 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6057 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6058 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6059 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6060 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6061 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6062 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6063 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6064 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6065 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6066 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6067 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6068 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6069 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6070 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6071 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6072 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6073 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6074 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6075 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6076 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6077 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

状态码： 202

表 5-6078 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6079 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((\text{.spec.completions} - \text{.status.successful}) < \text{.spec.parallelism})$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	参数类型	描述
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-6080 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6081 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "[topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br).[topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br).svc.[topic/body/section/table/tgroup/tbody/row/entry/p/br {''}] (br)". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.

参数	参数类型	描述
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6082 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6083 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6084 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6085 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6086 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6087 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6088 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6089 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6090 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6091 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6092 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6093 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6094 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6095 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6096 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6097 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6098 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6099 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6100 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6101 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6102 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6103 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6104 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6105 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6106 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6107 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6108 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6109 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6110 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6111 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6112 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-6113 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-6114 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-6115 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-6116 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6117 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6118 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6119 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6120 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6121 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6122 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6123 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6124 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6125 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6126 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6127 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6128 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6129 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6130 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6131 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6132 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6133 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6134 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6135 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6136 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6137 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6138 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6139 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6140 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6141 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6142 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6143 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6144 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6145 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6146 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6147 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6148 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6149 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6150 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6151 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6152 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6153 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6154 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br):[/topic/body/section/table/tgroup/tbody/row/entry/p/br {""]} (br) will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.

参数	参数类型	描述
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6155 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <pre> [/<topic ""="" "-"="" (br)="" (br).="" (br).)="" (br).[="" (br)[="" (international="" (note="" .="" ...="" 1="" 9="" ::="Ki" <a="" <topic="" [="" be="" body="" br="" case="" ei="" empty,="" entry="" from="" gi="" href="http://physics.nist.gov/cuu/Units/binary.html" in="" may="" mi="" of="" p="" pi="" row="" section="" see:="" system="" table="" tbody="" tgroup="" that="" the="" ti="" units;="" {""}="" ="">http://physics.nist.gov/cuu/Units/binary.html) [/<topic <="" body="" pre="" section="" table="" tgroup=""> </topic></topic></pre>

参数	参数类型	描述
		<p>tbody/row/entry/p/br {""}) (br) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) ::= "e" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br) "E" [/topic/body/section/table/tgroup/tbody/row/entry/p/br {""}) (br)</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than $2^{63}-1$ in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that:</p> <ol style="list-style-type: none"> No precision is lost No fractional digits will be emitted The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative. <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-6156 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6157 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6158 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6159 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6160 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6161 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project

参数	参数类型	描述
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6162 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6163 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6164 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6165 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6166 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6167 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6168 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6169 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6170 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6171 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6172 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6173 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6174 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6175 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6176 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6177 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

请求示例

- 创建普通Job，计算 π 到2000位并打印输出。

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "name": "pi"
  },
  "spec": {
    "template": {
      "metadata": {
        "name": "pi"
      },
      "spec": {
        "containers": [ {
          "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
          "image": "perl",
          "name": "pi",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1024Mi"
            },
            "requests": {
              "cpu": "500m",
              "memory": "1024Mi"
            }
          }
        }
      ],
      "imagePullSecrets": [ {
        "name": "imagepull-secret"
      } ],
      "priority": 0,
      "restartPolicy": "Never"
    }
  }
}
```

- 创建使用GPU的Job，gpu-418.126版本显示驱动。

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "annotations": {
      "cri.cci.io/gpu-driver": "gpu-418.126",
      "description": ""
    }
  },
}
```

```
"labels": { },
"name": "gpu-job-test",
"namespace": "cci-namespace-44173581"
},
"spec": {
  "template": {
    "metadata": {
      "annotations": {
        "cri.cci.io/gpu-driver": "gpu-418.126"
      },
      "name": "gpu-job-test"
    },
    "spec": {
      "containers": [ {
        "command": [ "sleep", "3600" ],
        "image": "library/nginx:latest",
        "lifecycle": { },
        "name": "container-0",
        "resources": {
          "limits": {
            "cpu": 4,
            "memory": "32Gi",
            "nvidia.com/gpu-tesla-v100-16GB": 1
          },
          "requests": {
            "cpu": 4,
            "memory": "32Gi",
            "nvidia.com/gpu-tesla-v100-16GB": 1
          }
        }
      } ],
      "imagePullSecrets": [ {
        "name": "imagepull-secret"
      } ],
      "restartPolicy": "Never"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2018-09-05T01:10:59Z",
    "labels": {
      "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
      "job-name": "pi"
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5391201",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",
    "uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
      }
    }
  }
}
```

```

    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
          "job-name": "pi"
        },
        "name": "pi"
      },
      "spec": {
        "containers": [ {
          "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
          "image": "perl",
          "imagePullPolicy": "Always",
          "name": "pi",
          "resources": {
            "limits": {
              "cpu": "500m",
              "memory": "1Gi"
            },
            "requests": {
              "cpu": "500m",
              "memory": "1Gi"
            }
          },
          "terminationMessagePath": "/dev/termination-log",
          "terminationMessagePolicy": "File"
        } ],
        "dnsPolicy": "ClusterFirst",
        "imagePullSecrets": [ {
          "name": "imagepull-secret"
        } ],
        "priority": 0,
        "restartPolicy": "Never",
        "schedulerName": "default-scheduler",
        "securityContext": { }
      }
    },
    "status": { }
  }
}

```

状态码

状态码	描述
200	OK
201	Created
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed

状态码	描述
406	NotAcceptable
409	AlreadyExists
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.5 删除 Job

功能介绍

删除Job。

调用方法

请参见[如何调用API](#)。

URI

DELETE /apis/batch/v1/namespaces/{namespace}/jobs/{name}

表 5-6178 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6179 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Integer	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	参数类型	描述
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6180 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-6181 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
dryRun	否	Array of strings	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	否	Long	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	是否必选	参数类型	描述
orphanDependents	否	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
preconditions	否	io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
propagationPolicy	否	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 5-6182 io.k8s.apimachinery.pkg.apis.meta.v1.Preconditions

参数	是否必选	参数类型	描述
resourceVersion	否	String	Specifies the target ResourceVersion
uid	否	String	Specifies the target UID.

响应参数

状态码： 200

表 5-6183 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6184 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6185 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"

参数	参数类型	描述
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-6186 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

状态码: 202

表 5-6187 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
message	String	A human-readable description of the status of this operation.

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure". More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6188 io.k8s.apimachinery.pkg.apis.meta.v1.StatusDetails

参数	参数类型	描述
causes	Array of io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action <ul style="list-style-type: none">for those errors this field may indicate how long to wait before taking the alternate action.

参数	参数类型	描述
uid	String	UID of the resource. (when there is a single resource which can be described). More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6189 io.k8s.apimachinery.pkg.apis.meta.v1.StatusCause

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 5-6190 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.

参数	参数类型	描述
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

- 只删除Job（对应Pod不删除）：

```
{
  "apiVersion": "v1",
  "gracePeriodSeconds": 0,
  "kind": "DeleteOptions",
  "propagationPolicy": "Orphan"
}
```

- 前台级联删除（按照Pod->Job的顺序进行删除）

```
{
  "apiVersion": "v1",
  "kind": "DeleteOptions",
  "propagationPolicy": "Foreground"
}
```

- 后台级联删除（按照Job->Pod的顺序进行删除）

```
{
  "apiVersion": "v1",
```



```
"kind": "DeleteOptions",  
"propagationPolicy": "Background"  
}
```

响应示例

状态码: 200

OK

```
{  
  "apiVersion": "batch/v1",  
  "kind": "Job",  
  "metadata": {  
    "creationTimestamp": "2018-09-05T01:10:59Z",  
    "deletionGracePeriodSeconds": 0,  
    "deletionTimestamp": "2018-09-05T03:27:34Z",  
    "finalizers": [ "orphan" ],  
    "labels": {  
      "app": "test2",  
      "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",  
      "job-name": "pi"  
    },  
    "name": "pi",  
    "namespace": "namespace-test",  
    "resourceVersion": "5415127",  
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",  
    "uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"  
  },  
  "spec": {  
    "backoffLimit": 6,  
    "completions": 1,  
    "parallelism": 1,  
    "selector": {  
      "matchLabels": {  
        "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"  
      }  
    },  
    "template": {  
      "metadata": {  
        "annotations": {  
          "cri.cci.io/container-type": "secure-container"  
        },  
        "creationTimestamp": null,  
        "labels": {  
          "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",  
          "job-name": "pi"  
        },  
        "name": "pi"  
      },  
      "spec": {  
        "containers": [ {  
          "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],  
          "image": "perl",  
          "imagePullPolicy": "Always",  
          "name": "pi",  
          "resources": {  
            "limits": {  
              "cpu": "500m",  
              "memory": "1Gi"  
            },  
            "requests": {  
              "cpu": "500m",  
              "memory": "1Gi"  
            }  
          }  
        },  
        "terminationMessagePath": "/dev/termination-log",  
        "terminationMessagePolicy": "File"  
      }  
    }  
  },  
}
```

```
"dnsPolicy": "ClusterFirst",
"imagePullSecrets": [ {
  "name": "imagepull-secret"
} ],
"restartPolicy": "Never",
"schedulerName": "default-scheduler",
"securityContext": { }
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}
```

状态码

状态码	描述
200	OK
202	Accepted
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.6 查询 Job

功能介绍

查询Job的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1/namespaces/{namespace}/jobs/{name}

表 5-6191 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6192 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6193 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-6194 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6195 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6

参数	参数类型	描述
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-6196 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6197 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6198 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6199 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6200 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6201 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6202 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6203 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6204 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6205 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6206 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6207 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6208 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6209 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6210 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6211 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6212 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6213 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6214 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6215 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6216 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6217 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6218 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6219 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6220 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6221 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6222 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6223 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6224 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6225 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6226 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6227 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6228 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6233 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6234 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6235 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6236 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6237 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6238 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6239 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6240 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6241 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6242 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6243 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6244 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6245 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6246 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6247 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6248 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6249 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6250 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6251 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6252 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6253 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6254 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6255 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6256 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6257 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6258 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6259 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6260 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6261 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6262 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6263 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6264 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6265 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6266 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6267 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6268 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6269 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6270 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6271 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6272 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6273 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6274 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6275 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6276 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6277 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6278 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6279 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6280 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6281 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6282 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6283 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6284 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6285 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6286 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6287 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6288 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6289 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6290 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6291 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6292 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6293 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2018-09-05T01:10:59Z",
    "labels": {
      "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4",
      "job-name": "pi"
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5391205",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",
    "uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4"
      }
    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4",
          "job-name": "pi"
        },
        "name": "pi"
      },
      "spec": {
```

```
"containers": [ {
  "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
  "image": "perl",
  "imagePullPolicy": "Always",
  "name": "pi",
  "resources": {
    "limits": {
      "cpu": "500m",
      "memory": "1Gi"
    },
    "requests": {
      "cpu": "500m",
      "memory": "1Gi"
    }
  },
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
}],
"dnsPolicy": "ClusterFirst",
"imagePullSecrets": [ {
  "name": "imagepull-secret"
}],
"restartPolicy": "Never",
"schedulerName": "default-scheduler",
"securityContext": { }
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.7 更新 Job

功能介绍

更新Job。

The following fields can be updated:

- metadata.labels
- spec.parallelism
- spec.completions
- spec.selector

调用方法

请参见[如何调用API](#)。

URI

PATCH /apis/batch/v1/namespaces/{namespace}/jobs/{name}

表 5-6294 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6295 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed

参数	是否必选	参数类型	描述
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint . This field is required for apply requests (application/apply-patch) but optional for non-apply patch types (JsonPatch, MergePatch, StrategicMergePatch).
force	否	Boolean	Force is going to "force" Apply requests. It means user will re-acquire conflicting fields owned by other people. Force flag must be unset for non-apply patch requests.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6296 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

参数	是否必选	参数类型	描述
Content-Type	是	String	<p>目前支持三种类型的PATCH请求方法的操作，参考《使用JSON合并patch更新Deployment》。</p> <ol style="list-style-type: none">1. Json Patch, Content-Type: application/json-patch+json 在RFC6902协议的定义中，Json Patch包含一系列对目标JSON对象的操作，其本身也为JSON对象。服务器接收到该对象后，会将其表示的操作应用于目标JSON对象。2. Merge Patch, Content-Type: application/merge-patch+json 在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。3. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

表 5-6297 请求 Body 参数

参数	是否必选	参数类型	描述
-	是	Object	Patch is provided to give a concrete name and type to the Kubernetes PATCH request body.

响应参数

状态码： 200

表 5-6298 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6299 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6

参数	参数类型	描述
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
<code>ttlSecondsAfterFinished</code>	Integer	<code>ttlSecondsAfterFinished</code> limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, <code>ttlSecondsAfterFinished</code> after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the <code>TTLAfterFinished</code> feature.

表 5-6300 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
<code>metadata</code>	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
<code>spec</code>	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6301 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
<code>activeDeadlineSeconds</code>	Long	Optional duration in seconds the pod may be active on the node relative to <code>StartTime</code> before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
<code>affinity</code>	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
<code>automountServiceAccountToken</code>	Boolean	<code>AutomountServiceAccountToken</code> indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6302 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6303 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6304 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6305 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6306 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6307 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6308 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6309 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6310 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6311 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6312 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6313 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6314 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6315 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6316 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6317 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6318 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6319 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6320 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6321 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6322 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6323 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6324 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6325 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6326 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6327 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6328 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6329 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6330 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6331 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6332 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-6333 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-6334 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-6335 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-6336 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6337 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6338 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6339 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6340 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6341 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6342 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6343 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6344 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6345 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6346 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6347 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	<p>FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.</p>
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	<p>Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running</p>
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	<p>GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk</p>

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6348 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6349 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6350 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6351 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6352 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6353 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6354 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6355 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6356 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6357 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6358 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6359 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6360 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6361 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6362 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6363 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6364 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6365 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6366 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6367 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6368 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6369 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6370 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6371 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6372 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6373 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6374 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6375 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6376 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6377 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6378 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6379 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6380 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6381 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6382 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6383 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6384 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6385 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6386 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6387 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6388 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6389 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6390 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6391 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6392 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6393 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6394 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6395 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6396 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6397 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

请求示例

更新Job中的labels值为"app:test"。

```
{
  "metadata": {
    "labels": {
      "app": "test"
    }
  }
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2018-09-05T01:10:59Z",
    "labels": {
      "app": "test",
      "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
      "job-name": "pi"
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5398083",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",
    "uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
      }
    }
  },
  "template": {
    "metadata": {
      "annotations": {
        "cri.cci.io/container-type": "secure-container"
      }
    }
  }
}
```

```
},
"creationTimestamp": null,
"labels": {
  "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
  "job-name": "pi"
},
"name": "pi"
},
"spec": {
  "containers": [ {
    "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
    "image": "perl",
    "imagePullPolicy": "Always",
    "name": "pi",
    "resources": {
      "limits": {
        "cpu": "500m",
        "memory": "1Gi"
      },
      "requests": {
        "cpu": "500m",
        "memory": "1Gi"
      }
    },
    "terminationMessagePath": "/dev/termination-log",
    "terminationMessagePolicy": "File"
  } ],
  "dnsPolicy": "ClusterFirst",
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "restartPolicy": "Never",
  "schedulerName": "default-scheduler",
  "securityContext": { }
}
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType

状态码	描述
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.20.8 替换 Job

功能介绍

替换Job。

The following fields can be updated:

- metadata.selfLink
- metadata.resourceVersion
- metadata.generation
- metadata.creationTimestamp
- metadata.deletionTimestamp
- metadata.labels
- metadata.clusterName
- metadata.generateName
- metadata.annotations
- spec.replicas
- template.containers
- spec.restartPolicy
- spec.activeDeadlineSeconds

调用方法

请参见[如何调用API](#)。

URI

PUT /apis/batch/v1/namespaces/{namespace}/jobs/{name}

表 5-6398 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6399 Query 参数

参数	是否必选	参数类型	描述
dryRun	否	String	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
fieldManager	否	String	fieldManager is a name associated with the actor or entity that is making these changes. The value must be less than or 128 characters long, and only contain printable characters, as defined by https://golang.org/pkg/unicode/#IsPrint .
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6400 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。
Content-Type	是	String	消息体的类型（格式），默认取值为“application/json” 缺省值： application/json

表 5-6401 请求 Body 参数

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	否	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	否	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6402 io.k8s.api.batch.v1.JobSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	否	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	否	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	是否必选	参数类型	描述
manualSelector	否	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	否	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	是否必选	参数类型	描述
template	是	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/job-run-to-completion/
tTlSecondsAfterFinished	否	Integer	tTlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, tTlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-6403 io.k8s.api.core.v1.PodTemplateSpec

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	否	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6404 io.k8s.api.core.v1.PodSpec

参数	是否必选	参数类型	描述
activeDeadlineSeconds	否	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	否	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	否	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	是	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	否	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	否	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	否	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.

参数	是否必选	参数类型	描述
ephemeralContainers	否	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	否	Array of io.k8s.api.core.v1.HostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	否	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	否	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	否	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	否	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.

参数	是否必选	参数类型	描述
imagePullSecrets	否	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	否	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/

参数	是否必选	参数类型	描述
nodeName	否	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	否	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	否	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.

参数	是否必选	参数类型	描述
preemptionPolicy	否	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	否	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	否	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	否	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	是否必选	参数类型	描述
restartPolicy	否	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	否	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	否	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	否	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceName	否	String	DeprecatedServiceAccount is a depreciated alias for ServiceAccountName. Deprecated: Use serviceName instead.

参数	是否必选	参数类型	描述
serviceAccountName	否	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	否	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	否	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	否	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.

参数	是否必选	参数类型	描述
terminationGracePeriodSeconds	否	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	否	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	否	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	否	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6405 io.k8s.api.core.v1.Affinity

参数	是否必选	参数类型	描述
nodeAffinity	否	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	是否必选	参数类型	描述
podAffinity	否	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	否	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6406 io.k8s.api.core.v1.NodeAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	否	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6407 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	是否必选	参数类型	描述
preference	是	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	是	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6408 io.k8s.api.core.v1.NodeSelectorTerm

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	否	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6409 io.k8s.api.core.v1.NodeSelector

参数	是否必选	参数类型	描述
nodeSelectorTerms	是	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6410 io.k8s.api.core.v1.NodeSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	The label key that the selector applies to.
operator	是	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.
values	否	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6411 io.k8s.api.core.v1.PodAffinity

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringSchedulingIgnoredDuringExecution</code> affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.

表 5-6412 io.k8s.api.core.v1.PodAntiAffinity

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	否	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6413 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	是否必选	参数类型	描述
podAffinityTerm	是	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	是	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6414 io.k8s.api.core.v1.PodAffinityTerm

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	否	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	是	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6415 io.k8s.api.core.v1.PodDNSConfig

参数	是否必选	参数类型	描述
nameservers	否	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	否	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	否	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6416 io.k8s.api.core.v1.PodDNSConfigOption

参数	是否必选	参数类型	描述
name	否	String	Required.
value	否	String	value is the value of the option

表 5-6417 io.k8s.api.core.v1.EphemeralContainer

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	是	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.

参数	是否必选	参数类型	描述
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	否	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	否	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
targetContainerName	否	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	是否必选	参数类型	描述
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6418 io.k8s.api.core.v1.HostAlias

参数	是否必选	参数类型	描述
hostnames	否	Array of strings	Hostnames for the above IP address.
ip	否	String	IP address of the host file entry.

表 5-6419 io.k8s.api.core.v1.LocalObjectReference

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6420 io.k8s.api.core.v1.Container

参数	是否必选	参数类型	描述
args	否	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	否	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$ \$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	否	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.

参数	是否必选	参数类型	描述
envFrom	否	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	否	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	否	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	否	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

参数	是否必选	参数类型	描述
name	是	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	否	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	否	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/management-compute-resources-container/
securityContext	否	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/

参数	是否必选	参数类型	描述
startupProbe	否	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	否	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	否	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false

参数	是否必选	参数类型	描述
terminationMessagePath	否	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	否	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	否	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	否	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	否	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.

参数	是否必选	参数类型	描述
workingDir	否	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6421 io.k8s.api.core.v1.EnvVar

参数	是否必选	参数类型	描述
name	是	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	否	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	否	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6422 io.k8s.api.core.v1.EnvVarSource

参数	是否必选	参数类型	描述
configMapKeyRef	否	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP, status.podIPs.
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	否	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6423 io.k8s.api.core.v1.ConfigMapKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key to select.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6424 io.k8s.api.core.v1.SecretKeySelector

参数	是否必选	参数类型	描述
key	是	String	The key of the secret to select from. Must be a valid secret key.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-6425 io.k8s.api.core.v1.EnvFromSource

参数	是否必选	参数类型	描述
configMapRef	否	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	否	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	否	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6426 io.k8s.api.core.v1.ConfigMapEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap must be defined

表 5-6427 io.k8s.api.core.v1.SecretEnvSource

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret must be defined

表 5-6428 io.k8s.api.core.v1.Lifecycle

参数	是否必选	参数类型	描述
postStart	否	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

参数	是否必选	参数类型	描述
preStop	否	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6429 io.k8s.api.core.v1.Handler

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6430 io.k8s.api.core.v1.ContainerPort

参数	是否必选	参数类型	描述
containerPort	是	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	否	String	What host IP to bind the external port to.
hostPort	否	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	否	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	否	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6431 io.k8s.api.core.v1.SecurityContext

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	否	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

参数	是否必选	参数类型	描述
capabilities	否	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	否	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.
procMount	否	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	否	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6432 io.k8s.api.core.v1.Capabilities

参数	是否必选	参数类型	描述
add	否	Array of strings	Added capabilities
drop	否	Array of strings	Removed capabilities

表 5-6433 io.k8s.api.core.v1.Probe

参数	是否必选	参数类型	描述
exec	否	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	否	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	否	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	否	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	否	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	否	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.
tcpSocket	否	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	否	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6434 io.k8s.api.core.v1.ExecAction

参数	是否必选	参数类型	描述
command	否	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6435 io.k8s.api.core.v1.HTTPGetAction

参数	是否必选	参数类型	描述
host	否	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	否	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	否	String	Path to access on the HTTP server.
port	是	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	否	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-6436 io.k8s.api.core.v1.HTTPHeader

参数	是否必选	参数类型	描述
name	是	String	The header field name

参数	是否必选	参数类型	描述
value	是	String	The header field value

表 5-6437 io.k8s.api.core.v1.TCPSocketAction

参数	是否必选	参数类型	描述
host	否	String	Optional: Host name to connect to, defaults to the pod IP.
port	是	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-6438 io.k8s.api.core.v1.VolumeDevice

参数	是否必选	参数类型	描述
devicePath	是	String	devicePath is the path inside of the container that the device will be mapped to.
name	是	String	name must match the name of a persistentVolumeClaim in the pod

表 5-6439 io.k8s.api.core.v1.VolumeMount

参数	是否必选	参数类型	描述
extendPathMode	否	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	是	String	Path within the container at which the volume should be mounted. Must not contain ':'.

参数	是否必选	参数类型	描述
mountPropagation	否	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	是	String	This must match the Name of a Volume.
policy	否	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	否	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	否	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	否	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$ (VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6440 io.k8s.api.core.v1.Policy

参数	是否必选	参数类型	描述
logs	否	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6441 io.k8s.api.core.v1.Logs

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations for log.
rotate	是	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6442 io.k8s.api.core.v1.PodReadinessGate

参数	是否必选	参数类型	描述
conditionType	是	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6443 io.k8s.api.core.v1.PodSecurityContext

参数	是否必选	参数类型	描述
fsGroup	否	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>

参数	是否必选	参数类型	描述
fsGroupChangePolicy	否	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	否	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	否	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

参数	是否必选	参数类型	描述
runAsUser	否	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seLinuxOptions	否	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	否	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	否	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	否	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	否	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6444 io.k8s.api.core.v1.SELinuxOptions

参数	是否必选	参数类型	描述
level	否	String	Level is SELinux level label that applies to the container.
role	否	String	Role is a SELinux role label that applies to the container.
type	否	String	Type is a SELinux type label that applies to the container.
user	否	String	User is a SELinux user label that applies to the container.

表 5-6445 io.k8s.api.core.v1.SeccompProfile

参数	是否必选	参数类型	描述
localhostProfile	否	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	是	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6446 io.k8s.api.core.v1.Sysctl

参数	是否必选	参数类型	描述
name	是	String	Name of a property to set
value	是	String	Value of a property to set

表 5-6447 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	是否必选	参数类型	描述
gmsaCredentialSpec	否	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the <code>GMSACredentialSpecName</code> field.
gmsaCredentialSpecName	否	String	<code>GMSACredentialSpecName</code> is the name of the GMSA credential spec to use.
runAsUserName	否	String	The <code>UserName</code> in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in <code>PodSecurityContext</code> . If set in both <code>SecurityContext</code> and <code>PodSecurityContext</code> , the value specified in <code>SecurityContext</code> takes precedence.

表 5-6448 io.k8s.api.core.v1.Toleration

参数	是否必选	参数类型	描述
effect	否	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are <code>NoSchedule</code> , <code>PreferNoSchedule</code> and <code>NoExecute</code> .
key	否	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be <code>Exists</code> ; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	否	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	否	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	否	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6449 io.k8s.api.core.v1.TopologySpreadConstraint

参数	是否必选	参数类型	描述
labelSelector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	是否必选	参数类型	描述
maxSkew	是	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1 (zone2) would make the ActualSkew(2-0) on zone1 (zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	是	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.

参数	是否必选	参数类型	描述
whenUnsatisfiable	是	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6450 io.k8s.api.core.v1.Volume

参数	是否必选	参数类型	描述
awsElasticBlockStore	否	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
azureDisk	否	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	否	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	否	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	否	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	否	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	否	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	否	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	是否必选	参数类型	描述
ephemeral	否	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	否	io.k8s.api.core.v1.FCVolumeSource object	<p>FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.</p>

参数	是否必选	参数类型	描述
flexVolume	否	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	否	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	否	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
gitRepo	否	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	否	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	否	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

参数	是否必选	参数类型	描述
iscsi	否	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	否	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	是	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	否	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	否	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	否	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine
portworxVolume	否	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	否	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API

参数	是否必选	参数类型	描述
quobyte	否	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	否	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	否	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	否	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	否	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	否	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6451 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

参数	是否必选	参数类型	描述
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	否	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	是	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6452 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	是否必选	参数类型	描述
cachingMode	否	String	Host Caching mode: None, Read Only, Read Write.
diskName	是	String	The Name of the data disk in the blob storage
diskURI	是	String	The URI the data disk in the blob storage
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.

参数	是否必选	参数类型	描述
kind	否	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6453 io.k8s.api.core.v1.AzureFileVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	是	String	the name of secret that contains Azure Storage Account Name and Key
shareName	是	String	Share Name

表 5-6454 io.k8s.api.core.v1.CephFSVolumeSource

参数	是否必选	参数类型	描述
monitors	是	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	否	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	否	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	否	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6455 io.k8s.api.core.v1.CinderVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

参数	是否必选	参数类型	描述
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	是	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6456 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	是否必选	参数类型	描述
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6457 io.k8s.api.core.v1.CSIVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	否	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	否	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeAttributes	否	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6458 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6459 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	是否必选	参数类型	描述
medium	否	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	否	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6460 io.k8s.api.core.v1.EphemeralVolumeSource

参数	是否必选	参数类型	描述
readOnly	否	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	是否必选	参数类型	描述
volumeClaimTemplate	否	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a standalone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <code><pod name>-<volume name></code> where <code><volume name></code> is the name from the <code>PodSpec.Volumes</code> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6461 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	是否必选	参数类型	描述
metadata	否	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	是	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6462 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	是否必选	参数类型	描述
annotations	否	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	否	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	是否必选	参数类型	描述
creationTimestamp	否	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	否	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	是否必选	参数类型	描述
deletionTimestamp	否	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	是否必选	参数类型	描述
finalizers	否	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.

参数	是否必选	参数类型	描述
generateName	否	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency</p>
generation	否	Long	<p>A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.</p>
labels	否	Map<String,String>	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels</p>

参数	是否必选	参数类型	描述
managedFields	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	否	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	否	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces

参数	是否必选	参数类型	描述
ownerReferences	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
resourceVersion	否	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	否	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

参数	是否必选	参数类型	描述
uid	否	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6463 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	是否必选	参数类型	描述
apiVersion	否	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	否	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	否	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.
manager	否	String	Manager is an identifier of the workflow managing these fields.
operation	否	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	否	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6464 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	是否必选	参数类型	描述
apiVersion	是	String	API version of the referent.
blockOwnerDeletion	否	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	否	Boolean	If true, this reference points to the managing controller.
kind	是	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	是	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	是	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6465 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	是否必选	参数类型	描述
accessModes	否	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	是否必选	参数类型	描述
dataSource	否	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	否	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	否	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	否	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1

参数	是否必选	参数类型	描述
volumeMode	否	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	否	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6466 io.k8s.api.core.v1.TypedLocalObjectReference

参数	是否必选	参数类型	描述
apiGroup	否	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	是	String	Kind is the type of resource being referenced
name	是	String	Name is the name of resource being referenced

表 5-6467 io.k8s.api.core.v1.ResourceRequirements

参数	是否必选	参数类型	描述
limits	否	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	是否必选	参数类型	描述
requests	否	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6468 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	是否必选	参数类型	描述
matchExpressions	否	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	否	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6469 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	是否必选	参数类型	描述
key	是	String	key is the label key that the selector applies to.
operator	是	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

参数	是否必选	参数类型	描述
values	否	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6470 io.k8s.api.core.v1.FCVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	否	Integer	Optional: FC target lun number
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	否	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	否	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6471 io.k8s.api.core.v1.FlexVolumeSource

参数	是否必选	参数类型	描述
driver	是	String	Driver is the name of the driver to use for this volume.

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.
options	否	Map<String,String>	Optional: Extra command options if any.
readOnly	否	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6472 io.k8s.api.core.v1.FlockerVolumeSource

参数	是否必选	参数类型	描述
datasetName	否	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	否	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6473 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	否	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
pdName	是	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6474 io.k8s.api.core.v1.GitRepoVolumeSource

参数	是否必选	参数类型	描述
directory	否	String	Target directory name. Must not contain or start with '!'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	是	String	Repository URL
revision	否	String	Commit hash for the specified revision.

表 5-6475 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	是否必选	参数类型	描述
endpoints	是	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	是	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	否	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6476 io.k8s.api.core.v1.HostPathVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	否	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6477 io.k8s.api.core.v1.ISCSIVolumeSource

参数	是否必选	参数类型	描述
chapAuthDiscovery	否	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	否	Boolean	whether support iSCSI Session CHAP authentication
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	否	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	是	String	Target iSCSI Qualified Name.
iscsiInterface	否	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	是	Integer	iSCSI Target Lun number.

参数	是否必选	参数类型	描述
portals	否	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication
targetPortal	是	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6478 io.k8s.api.core.v1.LocalDirVolumeSource

参数	是否必选	参数类型	描述
sizeLimit	否	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or</p>

参数	是否必选	参数类型	描述
			<p>suffix) is as large as possible. The sign will be omitted unless the number is negative. Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that will cause implementors to also use a fixed point implementation.</p>

表 5-6479 io.k8s.api.core.v1.NFSVolumeSource

参数	是否必选	参数类型	描述
path	是	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	否	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	是	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6480 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	是否必选	参数类型	描述
claimName	是	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	否	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6481 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	是	String	ID that identifies Photon Controller persistent disk

表 5-6482 io.k8s.api.core.v1.PortworxVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	FSType represents the filesystem type to mount Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	是	String	VolumeID uniquely identifies a Portworx volume

表 5-6483 io.k8s.api.core.v1.ProjectedVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	是	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6484 io.k8s.api.core.v1.VolumeProjection

参数	是否必选	参数类型	描述
configMap	否	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	否	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project
secret	否	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	否	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6485 io.k8s.api.core.v1.ConfigMapProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6486 io.k8s.api.core.v1.DownwardAPIProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6487 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	是否必选	参数类型	描述
fieldRef	否	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.

参数	是否必选	参数类型	描述
mode	否	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	是	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	否	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6488 io.k8s.api.core.v1.ObjectFieldSelector

参数	是否必选	参数类型	描述
apiVersion	否	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	是	String	Path of the field to select in the specified API version.

表 5-6489 io.k8s.api.core.v1.ResourceFieldSelector

参数	是否必选	参数类型	描述
containerName	否	String	Container name: required for volumes, optional for env vars

参数	是否必选	参数类型	描述
divisor	否	String	Specifies the output format of the exposed resources, defaults to "1"
resource	是	String	Required: resource to select

表 5-6490 io.k8s.api.core.v1.SecretProjection

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	否	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	否	Boolean	Specify whether the Secret or its key must be defined

表 5-6491 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	是否必选	参数类型	描述
audience	否	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	否	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	是	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6492 io.k8s.api.core.v1.QuobyteVolumeSource

参数	是否必选	参数类型	描述
group	否	String	Group to map volume access to Default is no group
readOnly	否	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	是	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes

参数	是否必选	参数类型	描述
tenant	否	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	否	String	User to map volume access to Defaults to serviceaccount user
volume	是	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6493 io.k8s.api.core.v1.RBDVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	是	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	否	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	是	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
pool	否	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	是否必选	参数类型	描述
readOnly	否	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	否	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6494 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	是	String	The host address of the ScaleIO API Gateway.
protectionDomain	否	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	是	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	否	Boolean	Flag to enable/disable SSL communication with Gateway, default false

参数	是否必选	参数类型	描述
storageMode	否	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	否	String	The ScaleIO Storage Pool associated with the protection domain.
system	是	String	The name of the storage system as configured in ScaleIO.
volumeName	否	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6495 io.k8s.api.core.v1.SecretVolumeSource

参数	是否必选	参数类型	描述
defaultMode	否	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
items	否	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	否	Boolean	Specify whether the Secret or its keys must be defined
secretName	否	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6496 io.k8s.api.core.v1.KeyToPath

参数	是否必选	参数类型	描述
key	是	String	The key to project.
mode	否	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

参数	是否必选	参数类型	描述
path	是	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6497 io.k8s.api.core.v1.StorageOSVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	否	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	否	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	否	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.

参数	是否必选	参数类型	描述
volumeNamespace	否	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6498 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	是否必选	参数类型	描述
fsType	否	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	否	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	否	String	Storage Policy Based Management (SPBM) profile name.
volumePath	是	String	Path that identifies vSphere volume vmk

表 5-6499 io.k8s.api.batch.v1.JobStatus

参数	是否必选	参数类型	描述
active	否	Integer	The number of actively running pods.

参数	是否必选	参数类型	描述
completionTime	否	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	否	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	否	Integer	The number of pods which reached phase Failed.
startTime	否	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	否	Integer	The number of pods which reached phase Succeeded.

表 5-6500 io.k8s.api.batch.v1.JobCondition

参数	是否必选	参数类型	描述
lastProbeTime	否	String	Last time the condition was checked.
lastTransitionTime	否	String	Last time the condition transit from one status to another.
message	否	String	Human readable message indicating details about last transition.
reason	否	String	(brief) reason for the condition's last transition.
status	是	String	Status of the condition, one of True, False, Unknown.
type	是	String	Type of job condition, Complete or Failed.

响应参数

状态码： 200

表 5-6501 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6502 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer

参数	参数类型	描述
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((.spec.completions - .status.successful) < .spec.parallelism)$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-6503 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6504 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6505 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6506 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6507 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6508 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6509 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6510 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6511 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6512 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6513 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6514 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6515 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6516 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6517 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6518 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6519 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6520 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6521 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6522 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6523 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6524 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6525 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6526 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6527 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6528 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6529 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6530 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6531 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6532 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6533 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6534 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6535 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-6536 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-6537 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-6538 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-6539 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '!'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6540 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6541 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6542 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6543 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6544 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6545 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6546 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6547 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entryptpoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6548 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6549 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6550 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6551 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6552 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6553 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6554 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6555 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6556 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6557 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6558 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6559 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6560 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6561 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6562 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6563 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6564 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6565 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6566 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6567 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6568 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6569 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6570 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6571 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6572 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6573 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6574 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6575 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6576 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6577 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6578 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6579 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6580 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6581 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6582 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6583 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6584 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6585 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6586 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6587 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6588 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6589 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6590 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6591 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6592 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6593 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6594 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6595 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6596 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6597 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6598 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6599 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6600 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

状态码： 201

表 5-6601 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6602 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $((\text{.spec.completions} - \text{.status.successful}) < \text{.spec.parallelism})$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors

参数	参数类型	描述
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
ttlSecondsAfterFinished	Integer	ttlSecondsAfterFinished limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, ttlSecondsAfterFinished after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the TTLAfterFinished feature.

表 5-6603 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6604 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints

参数	参数类型	描述
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.

参数	参数类型	描述
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

参数	参数类型	描述
readinessGate s	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClass Name	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerNa me	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityConte xt	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccoun t	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccoun tName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/

参数	参数类型	描述
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.

参数	参数类型	描述
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6605 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6606 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6607 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6608 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6609 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6610 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6611 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6612 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringSchedulingIgnoredDuringExecution anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6613 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6614 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6615 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.

参数	参数类型	描述
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6616 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6617 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell

参数	参数类型	描述
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6618 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6619 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6620 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6621 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6622 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <i>metadata.name</i> , <i>metadata.namespace</i> , <i>metadata.labels['<KEY>']</i> , <i>metadata.annotations['<KEY>']</i> , <i>spec.nodeName</i> , <i>spec.serviceAccountName</i> , <i>status.hostIP</i> , <i>status.podIP</i> , <i>status.podIPs</i> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6623 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6624 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6625 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6626 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6627 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6628 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6629 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6630 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6631 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6632 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6633 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6634 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6635 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 5-6636 io.k8s.api.core.v1.HTTPHeader

参数	参数类型	描述
name	String	The header field name
value	String	The header field value

表 5-6637 io.k8s.api.core.v1.TCPSocketAction

参数	参数类型	描述
host	String	Optional: Host name to connect to, defaults to the pod IP.
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 5-6638 io.k8s.api.core.v1.VolumeDevice

参数	参数类型	描述
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.
name	String	name must match the name of a persistentVolumeClaim in the pod

表 5-6639 io.k8s.api.core.v1.VolumeMount

参数	参数类型	描述
extendPathMode	String	Extend the volume path by appending the pod metadata to the path according to specified pattern. which provide a way of directory isolation and help prevent the writing conflict between different pods.
mountPath	String	Path within the container at which the volume should be mounted. Must not contain '!'. Must not contain '..'.
mountPropagation	String	mountPropagation determines how mounts are propagated from the host to container and the other way around. When not set, MountPropagationNone is used. This field is beta in 1.10.
name	String	This must match the Name of a Volume.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6640 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6641 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6642 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6643 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6644 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6645 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6646 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6647 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6648 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6649 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6650 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6651 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6652 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6653 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6654 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6655 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6656 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6657 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6658 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6659 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6660 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6661 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6662 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6663 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6664 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6665 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6666 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6667 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6668 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6669 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6670 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6671 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6672 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6673 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6674 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6675 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6676 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6677 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6678 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6679 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6680 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6681 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6682 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6683 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6684 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6685 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6686 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6687 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6688 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6689 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6690 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6691 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6692 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6693 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6694 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6695 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6696 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6697 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6698 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6699 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6700 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

请求示例

给已创建Job中增加一个labels值"app": "test2"。

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2018-09-05T01:10:59Z",
    "labels": {
      "app": "test2",
      "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4",
      "job-name": "pi"
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5398083",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",
    "uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4"
      }
    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4",
          "job-name": "pi"
        },
        "name": "pi"
      },
      "spec": {
        "containers": [ {
          "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
          "image": "perl",
          "imagePullPolicy": "Always",
          "name": "pi",
          "resources": {
            "limits": {
```



```
    "cpu": "500m",
    "memory": "1Gi"
  },
  "requests": {
    "cpu": "500m",
    "memory": "1Gi"
  }
},
"terminationMessagePath": "/dev/termination-log",
"terminationMessagePolicy": "File"
}],
"dnsPolicy": "ClusterFirst",
"imagePullSecrets": [ {
  "name": "imagepull-secret"
}],
"restartPolicy": "Never",
"schedulerName": "default-scheduler",
"securityContext": { }
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}
```

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2018-09-05T01:10:59Z",
    "labels": {
      "app": "test2",
      "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
      "job-name": "pi"
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5400771",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi",
    "uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4"
      }
    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "controller-uid": "8c923079-b0a8-11e8-8bcb-f898ef6c78b4",
          "job-name": "pi"
        },
        "name": "pi"
      }
    }
  }
}
```

```
},
"spec": {
  "containers": [ {
    "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
    "image": "perl",
    "imagePullPolicy": "Always",
    "name": "pi",
    "resources": {
      "limits": {
        "cpu": "500m",
        "memory": "1Gi"
      },
      "requests": {
        "cpu": "500m",
        "memory": "1Gi"
      }
    },
    "terminationMessagePath": "/dev/termination-log",
    "terminationMessagePolicy": "File"
  } ],
  "dnsPolicy": "ClusterFirst",
  "imagePullSecrets": [ {
    "name": "imagepull-secret"
  } ],
  "restartPolicy": "Never",
  "schedulerName": "default-scheduler",
  "securityContext": { }
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}
```

状态码

状态码	描述
200	OK
201	Created
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError

状态码	描述
503	ServiceUnavailable
504	ServerTimeout

5.20.9 查询 Job 状态

功能介绍

查询Job状态。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status

表 5-6701 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the Job
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6702 Query 参数

参数	是否必选	参数类型	描述
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6703 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-6704 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.batch.v1.JobSpec object	Specification of the desired behavior of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.batch.v1.JobStatus object	Current status of a job. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6705 io.k8s.api.batch.v1.JobSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	Integer	Specifies the number of retries before marking this job failed. Defaults to 6

参数	参数类型	描述
completions	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
manualSelector or	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave <i>manualSelector</i> unset unless you are certain what you are doing. When false or unset, the system pick labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see <i>manualSelector=true</i> in jobs that were created with the old <i>extensions/v1beta1</i> API. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/#specifying-your-own-pod-selector
parallelism	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$, i.e. when the work left to do is less than max parallelism. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over pods that should match the pod count. Normally, the system sets this field for you. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Describes the pod that will be created when executing a job. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/

参数	参数类型	描述
<code>ttlSecondsAfterFinished</code>	Integer	<code>ttlSecondsAfterFinished</code> limits the lifetime of a Job that has finished execution (either Complete or Failed). If this field is set, <code>ttlSecondsAfterFinished</code> after the Job finishes, it is eligible to be automatically deleted. When the Job is being deleted, its lifecycle guarantees (e.g. finalizers) will be honored. If this field is unset, the Job won't be automatically deleted. If this field is set to zero, the Job becomes eligible to be deleted immediately after it finishes. This field is alpha-level and is only honored by servers that enable the <code>TTLAfterFinished</code> feature.

表 5-6706 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
<code>metadata</code>	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
<code>spec</code>	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6707 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
<code>activeDeadlineSeconds</code>	Long	Optional duration in seconds the pod may be active on the node relative to <code>StartTime</code> before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
<code>affinity</code>	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
<code>automountServiceAccountToken</code>	Boolean	<code>AutomountServiceAccountToken</code> indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6708 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6709 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6710 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6711 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6712 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6713 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6714 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6715 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6716 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6717 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6718 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6719 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6720 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6721 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6722 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6723 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manager-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6724 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6725 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6726 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6727 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6728 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6729 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6730 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6731 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6732 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6733 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6734 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6735 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6736 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6737 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6738 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6743 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6744 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6745 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6746 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6747 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6748 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6749 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6750 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entryptpoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6751 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6752 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6753 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6754 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6755 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6756 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6757 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6758 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6759 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6760 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6761 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6762 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6763 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6764 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6765 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6766 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.
fieldsType	String	FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"
fieldsV1	Object	FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6767 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6768 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6769 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6770 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6771 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6772 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6773 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6774 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6775 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6776 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6777 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6778 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6779 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6780 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6781 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6782 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6783 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6784 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6785 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6786 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6787 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6788 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6789 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6790 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6791 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6792 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6793 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6794 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6795 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6796 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6797 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6798 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6799 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6800 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6801 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6802 io.k8s.api.batch.v1.JobStatus

参数	参数类型	描述
active	Integer	The number of actively running pods.
completionTime	String	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
conditions	Array of io.k8s.api.batch.v1.JobCondition objects	The latest available observations of an object's current state. More info: https://kubernetes.io/docs/concepts/workloads/controllers/jobs-run-to-completion/
failed	Integer	The number of pods which reached phase Failed.
startTime	String	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	Integer	The number of pods which reached phase Succeeded.

表 5-6803 io.k8s.api.batch.v1.JobCondition

参数	参数类型	描述
lastProbeTime	String	Last time the condition was checked.

参数	参数类型	描述
lastTransition Time	String	Last time the condition transit from one status to another.
message	String	Human readable message indicating details about last transition.
reason	String	(brief) reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of job condition, Complete or Failed.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "batch/v1",
  "kind": "Job",
  "metadata": {
    "creationTimestamp": "2018-09-05T01:10:59Z",
    "labels": {
      "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4",
      "job-name": "pi"
    },
    "name": "pi",
    "namespace": "namespace-test",
    "resourceVersion": "5391205",
    "selfLink": "/apis/batch/v1/namespaces/namespace-test/jobs/pi/status",
    "uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4"
  },
  "spec": {
    "backoffLimit": 6,
    "completions": 1,
    "parallelism": 1,
    "selector": {
      "matchLabels": {
        "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4"
      }
    },
    "template": {
      "metadata": {
        "annotations": {
          "cri.cci.io/container-type": "secure-container"
        },
        "creationTimestamp": null,
        "labels": {
          "controller-uid": "8c923079-b0a8-11e8-8bcf-f898ef6c78b4",
          "job-name": "pi"
        },
        "name": "pi"
      },
      "spec": {
```

```
"containers": [ {
  "command": [ "perl", "-Mbignum=bpi", "-wle", "print bpi(2000)" ],
  "image": "perl",
  "imagePullPolicy": "Always",
  "name": "pi",
  "resources": {
    "limits": {
      "cpu": "500m",
      "memory": "1Gi"
    },
    "requests": {
      "cpu": "500m",
      "memory": "1Gi"
    }
  },
  "terminationMessagePath": "/dev/termination-log",
  "terminationMessagePolicy": "File"
}],
"dnsPolicy": "ClusterFirst",
"imagePullSecrets": [ {
  "name": "imagepull-secret"
}],
"restartPolicy": "Never",
"schedulerName": "default-scheduler",
"securityContext": { }
}
},
"status": {
  "active": 1,
  "startTime": "2018-09-05T01:10:59Z"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.21 ReplicaSet

5.21.1 查询指定 namespace 下的 ReplicaSets

功能介绍

查询命名空间下所有的ReplicaSets。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/replicasets

表 5-6804 路径参数

参数	是否必选	参数类型	描述
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6805 Query 参数

参数	是否必选	参数类型	描述
allowWatchBookmarks	否	Boolean	allowWatchBookmarks requests watch events with type "BOOKMARK". Servers that do not implement bookmarks may ignore this flag and bookmarks are sent at the server's discretion. Clients should not assume bookmarks are returned at any specific interval, nor may they assume the server will send any BOOKMARK event during a session. If this is not a watch, this field is ignored. If the feature gate WatchBookmarks is not enabled in apiserver, this field is ignored.

参数	是否必选	参数类型	描述
continue	否	String	<p>The continue option should be set when retrieving more results from the server. Since this value is server defined, clients may only use the continue value from a previous query result with identical query parameters (except for the value of continue) and the server may reject a continue value it does not recognize. If the specified continue value is no longer valid whether due to expiration (generally five to fifteen minutes) or a configuration change on the server, the server will respond with a 410 ResourceExpired error together with a continue token. If the client needs a consistent list, it must restart their list without the continue field. Otherwise, the client may send another list request with the token received with the 410 error, the server will respond with a list starting from the next key, but from the latest snapshot, which is inconsistent from the previous list results - objects that are created, modified, or deleted after the first list request will be included in the response, as long as their keys are after the "next key".</p> <p>This field is not supported when watch is true. Clients may start a watch from the last resourceVersion value returned by the server and not miss any modifications.</p>
fieldSelector	否	String	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	否	String	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	参数类型	描述
limit	否	Integer	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the <i>continue</i> field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	参数类型	描述
resourceVersion	否	String	resourceVersion sets a constraint on what resource versions a request may be served from. See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
resourceVersionMatch	否	String	resourceVersionMatch determines how resourceVersion is applied to list calls. It is highly recommended that resourceVersionMatch be set for list calls where resourceVersion is set See https://kubernetes.io/docs/reference/using-api/api-concepts/#resource-versions for details. Defaults to unset
timeoutSeconds	否	Integer	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	否	Boolean	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6806 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-6807 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
items	Array of io.k8s.api.apps.v1.ReplicaSet objects	List of ReplicaSets. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta object	Standard list metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

表 5-6808 io.k8s.api.apps.v1.ReplicaSet

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	If the Labels of a ReplicaSet are empty, they are defaulted to be the same as the Pod(s) that the ReplicaSet manages. Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.apps.v1.ReplicaSetSpec object	Spec defines the specification of the desired behavior of the ReplicaSet. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.apps.v1.ReplicaSetStatus object	Status is the most recently observed status of the ReplicaSet. This data may be out of date by some window of time. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6809 io.k8s.api.apps.v1.ReplicaSetSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
replicas	Integer	Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Selector is a label query over pods that should match the replica count. Label keys and values that must match in order to be controlled by this replica set. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Template is the object that describes the pod that will be created if insufficient replicas are detected. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#pod-template

表 5-6810 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6811 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.

参数	参数类型	描述
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostname	String	Specifies the hostname of the Pod If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images#specifying-imagepullsecrets-on-a-pod

参数	参数类型	描述
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	nodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6812 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6813 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6814 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6815 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6816 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6817 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6818 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6819 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6820 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6821 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6822 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6823 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6824 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6825 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6826 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6827 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6828 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6829 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6830 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6831 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6832 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6833 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6834 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6835 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6836 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6837 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6838 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6839 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6840 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6841 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6842 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6847 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6848 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6849 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6850 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- If unset, the Kubelet will not modify the ownership and permissions of any volume.
fsGroupChangePolicy	String	fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6851 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6852 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6853 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6854 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6855 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6856 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6857 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6858 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6859 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6860 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6861 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6862 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6863 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6864 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6865 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6866 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6867 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6868 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6869 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6870 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6871 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6872 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6873 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6874 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6875 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6876 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6877 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6878 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6879 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6880 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6881 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6882 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6883 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6884 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6885 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6886 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6887 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6888 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6889 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6890 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6891 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6892 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6893 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6894 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6895 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-6896 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-6897 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6898 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-6899 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-6900 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-6901 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-6902 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-6903 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-6904 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-6905 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-6906 io.k8s.api.apps.v1.ReplicaSetStatus

参数	参数类型	描述
availableReplicas	Integer	The number of available replicas (ready for at least minReadySeconds) for this replica set.
conditions	Array of io.k8s.api.apps.v1.ReplicaSetCondition objects	Represents the latest available observations of a replica set's current state.
fullyLabeledReplicas	Integer	The number of pods that have labels matching the labels of the pod template of the replicaset.
observedGeneration	Long	ObservedGeneration reflects the generation of the most recently observed ReplicaSet.
readyReplicas	Integer	The number of ready replicas for this replica set.
replicas	Integer	Replicas is the most recently observed number of replicas. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller

表 5-6907 io.k8s.api.apps.v1.ReplicaSetCondition

参数	参数类型	描述
lastTransitionTime	String	The last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of replica set condition.

表 5-6908 io.k8s.apimachinery.pkg.apis.meta.v1.ListMeta

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a consistent list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response, unless you have received this token from an error message.
remainingItemCount	Long	remainingItemCount is the number of subsequent items in the list which are not included in this list response. If the list request contained label or field selectors, then the number of remaining items is unknown and the field will be left unset and omitted during serialization. If the list is complete (either because it is not chunking or because this is the last chunk), then there are no more remaining items and this field will be left unset and omitted during serialization. Servers older than v1.15 do not set this field. The intended use of the remainingItemCount is <i>estimating</i> the size of a collection. Clients should not rely on the remainingItemCount to be set or to be exact.

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency
selfLink	String	selfLink is a URL representing this object. Populated by the system. Read-only. DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.

请求示例

无

响应示例

状态码: 200

OK

```
{
  "apiVersion": "apps/v1",
  "items": [ {
    "metadata": {
      "annotations": {
        "deployment.kubernetes.io/desired-replicas": "1",
        "deployment.kubernetes.io/max-replicas": "2",
        "deployment.kubernetes.io/revision": "1"
      },
      "creationTimestamp": "2017-12-12T11:27:15Z",
      "generation": 1,
      "labels": {
        "name": "deployment-test",
        "pod-template-hash": "800400086"
      },
      "name": "deployment-test-800400086",
      "namespace": "default",
      "ownerReferences": [ {
        "apiVersion": "extensions/v1beta1",
        "blockOwnerDeletion": true,
        "controller": true,
        "kind": "Deployment",
        "name": "deployment-test",
        "uid": "6776d16b-df2f-11e7-961f-fa163ed139d5"
      } ],
      "resourceVersion": "784294",
      "selfLink": "/apis/extensions/v1beta1/namespaces/default/replicasets/deployment-test-800400086",
      "uid": "67775454-df2f-11e7-961f-fa163ed139d5"
    },
    "spec": {
```

```
"replicas" : 1,
"selector" : {
  "matchLabels" : {
    "name" : "deployment-test",
    "pod-template-hash" : "800400086"
  }
},
"template" : {
  "metadata" : {
    "creationTimestamp" : null,
    "labels" : {
      "name" : "deployment-test",
      "pod-template-hash" : "800400086"
    }
  },
  "spec" : {
    "containers" : [ {
      "image" : "172.16.5.235:20202/test/testnginx:v3",
      "imagePullPolicy" : "IfNotPresent",
      "name" : "deployment-test",
      "resources" : { },
      "terminationMessagePath" : "/dev/termination-log",
      "terminationMessagePolicy" : "File"
    } ],
    "dnsPolicy" : "ClusterFirst",
    "restartPolicy" : "Always",
    "schedulerName" : "default-scheduler",
    "securityContext" : { }
  }
},
"status" : {
  "fullyLabeledReplicas" : 1,
  "observedGeneration" : 1,
  "replicas" : 1
}, {
  "metadata" : {
    "annotations" : {
      "deployment.kubernetes.io/desired-replicas" : "2",
      "deployment.kubernetes.io/max-replicas" : "2",
      "deployment.kubernetes.io/revision" : "1"
    },
    "creationTimestamp" : "2017-12-13T07:28:36Z",
    "generation" : 1,
    "labels" : {
      "app" : "test-pv",
      "pod-template-hash" : "446136006"
    },
    "name" : "test-pv-446136006",
    "namespace" : "default",
    "ownerReferences" : [ {
      "apiVersion" : "extensions/v1beta1",
      "blockOwnerDeletion" : true,
      "controller" : true,
      "kind" : "Deployment",
      "name" : "test-pv",
      "uid" : "3b3d3a22-dfd7-11e7-961f-fa163ed139d5"
    } ],
    "resourceVersion" : "784441",
    "selfLink" : "/apis/extensions/v1beta1/namespaces/default/replicasets/test-pv-446136006",
    "uid" : "3b3dee35-dfd7-11e7-961f-fa163ed139d5"
  },
  "spec" : {
    "replicas" : 2,
    "selector" : {
      "matchLabels" : {
        "app" : "test-pv",
        "pod-template-hash" : "446136006"
      }
    }
  }
}
```

```
}
},
"template": {
  "metadata": {
    "annotations": {
      "metrics.alpha.kubernetes.io/custom-endpoints": "[{api:',path:',port:',names:']]"
    },
    "creationTimestamp": null,
    "labels": {
      "app": "test-pv",
      "pod-template-hash": "446136006"
    }
  },
  "spec": {
    "affinity": { },
    "containers": [ {
      "image": "172.16.5.235:20202/test/nginx:latest",
      "imagePullPolicy": "Always",
      "lifecycle": { },
      "name": "container-0",
      "resources": { },
      "terminationMessagePath": "/dev/termination-log",
      "terminationMessagePolicy": "File",
      "volumeMounts": [ {
        "mountPath": "/tmp0",
        "name": "www",
        "readOnly": true
      } ]
    } ],
    "dnsPolicy": "ClusterFirst",
    "imagePullSecrets": [ {
      "name": "imagepull-secret"
    } ],
    "restartPolicy": "Always",
    "schedulerName": "default-scheduler",
    "securityContext": { },
    "volumes": [ {
      "name": "www",
      "persistentVolumeClaim": {
        "claimName": "pvc1513149915389"
      }
    } ]
  }
}
},
"status": {
  "availableReplicas": 2,
  "fullyLabeledReplicas": 2,
  "observedGeneration": 1,
  "readyReplicas": 2,
  "replicas": 2
}
},
"kind": "ReplicaSetList",
"metadata": {
  "resourceVersion": "793174",
  "selfLink": "/apis/apps/v1/namespaces/default/replicasets"
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest

状态码	描述
401	Unauthorized
403	Forbidden
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

5.21.2 查询 ReplicaSet

功能介绍

查询ReplicaSet的详细信息。

调用方法

请参见[如何调用API](#)。

URI

GET /apis/apps/v1/namespaces/{namespace}/replicasets/{name}

表 5-6909 路径参数

参数	是否必选	参数类型	描述
name	是	String	name of the ReplicaSet
namespace	是	String	object name and auth scope, such as for teams and projects

表 5-6910 Query 参数

参数	是否必选	参数类型	描述
exact	否	Boolean	Should the export be exact. Exact export maintains cluster-specific fields like 'Namespace'. Deprecated. Planned for removal in 1.18.
export	否	Boolean	Should this value be exported. Export strips fields that a user can not specify. Deprecated. Planned for removal in 1.18.
pretty	否	String	If 'true', then the output is pretty printed.

请求参数

表 5-6911 请求 Header 参数

参数	是否必选	参数类型	描述
X-Auth-Token	是	String	用户Token。 通过调用IAM服务获取用户Token接口获取（响应消息头中X-Subject-Token的值）。

响应参数

状态码： 200

表 5-6912 响应 Body 参数

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#resources

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	If the Labels of a ReplicaSet are empty, they are defaulted to be the same as the Pod(s) that the ReplicaSet manages. Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.apps.v1.ReplicaSetSpec object	Spec defines the specification of the desired behavior of the ReplicaSet. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status
status	io.k8s.api.apps.v1.ReplicaSetStatus object	Status is the most recently observed status of the ReplicaSet. This data may be out of date by some window of time. Populated by the system. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6913 io.k8s.api.apps.v1.ReplicaSetSpec

参数	参数类型	描述
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
replicas	Integer	Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller

参数	参数类型	描述
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	Selector is a label query over pods that should match the replica count. Label keys and values that must match in order to be controlled by this replica set. It must match the pod template's labels. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/labels/#label-selectors
template	io.k8s.api.core.v1.PodTemplateSpec object	Template is the object that describes the pod that will be created if insufficient replicas are detected. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller#pod-template

表 5-6914 io.k8s.api.core.v1.PodTemplateSpec

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
spec	io.k8s.api.core.v1.PodSpec object	Specification of the desired behavior of the pod. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#spec-and-status

表 5-6915 io.k8s.api.core.v1.PodSpec

参数	参数类型	描述
activeDeadlineSeconds	Long	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
affinity	io.k8s.api.core.v1.Affinity object	If specified, the pod's scheduling constraints
automountServiceAccountToken	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.

参数	参数类型	描述
containers	Array of io.k8s.api.core.v1.Container objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
dnsConfig	io.k8s.api.core.v1.PodDNSConfig object	Specifies the DNS parameters of a pod. Parameters specified here will be merged to the generated DNS configuration based on DNSPolicy.
dnsPolicy	String	Set DNS policy for the pod. Defaults to "ClusterFirst". Valid values are 'ClusterFirstWithHostNet', 'ClusterFirst', 'Default' or 'None'. DNS parameters given in DNSConfig will be merged with the policy selected with DNSPolicy. To have DNS options set along with hostNetwork, you have to specify DNS policy explicitly to 'ClusterFirstWithHostNet'.
enableServiceLinks	Boolean	EnableServiceLinks indicates whether information about services should be injected into pod's environment variables, matching the syntax of Docker links. Optional: Defaults to true.
ephemeralContainers	Array of io.k8s.api.core.v1.EphemeralContainer objects	List of ephemeral containers run in this pod. Ephemeral containers may be run in an existing pod to perform user-initiated actions such as debugging. This list cannot be specified when creating a pod, and it cannot be modified by updating the pod spec. In order to add an ephemeral container to an existing pod, use the pod's ephemeralcontainers subresource. This field is alpha-level and is only honored by servers that enable the EphemeralContainers feature.
hostAliases	Array of io.k8s.api.core.v1.HostAlias objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.

参数	参数类型	描述
hostname	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
imagePullSecrets	Array of io.k8s.api.core.v1.LocalObjectReference objects	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. For example, in the case of docker, only DockerConfig type secrets are honored. More info: https://kubernetes.io/docs/concepts/containers/images/#specifying-imagepullsecrets-on-a-pod
initContainers	Array of io.k8s.api.core.v1.Container objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, Liveness probes, or Startup probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/init-containers/
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
nodeSelector	Map<String,String>	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. More info: https://kubernetes.io/docs/concepts/configuration/assign-pod-node/

参数	参数类型	描述
overhead	Map<String,String>	Overhead represents the resource overhead associated with running a pod for a given RuntimeClass. This field will be autopopulated at admission time by the RuntimeClass admission controller. If the RuntimeClass admission controller is enabled, overhead must not be set in Pod create requests. The RuntimeClass admission controller will reject Pod create requests which have the overhead already set. If RuntimeClass is configured and selected in the PodSpec, Overhead will be set to the value defined in the corresponding RuntimeClass, otherwise it will remain unset and treated as zero. More info: https://git.k8s.io/enhancements/keps/sig-node/20190226-pod-overhead.md This field is alpha-level as of Kubernetes v1.16, and is only honored by servers that enable the PodOverhead feature.
preemptionPolicy	String	PreemptionPolicy is the Policy for preempting pods with lower priority. One of Never, PreemptLowerPriority. Defaults to PreemptLowerPriority if unset. This field is beta-level, gated by the NonPreemptingPriority feature-gate.
priority	Integer	The priority value. Various system components use this field to find the priority of the pod. When Priority Admission Controller is enabled, it prevents users from setting this field. The admission controller populates this field from PriorityClassName. The higher the value, the higher the priority.
priorityClassName	String	If specified, indicates the pod's priority. "system-node-critical" and "system-cluster-critical" are two special keywords which indicate the highest priorities with the former being the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.
readinessGates	Array of io.k8s.api.core.v1.PodReadinessGate objects	If specified, all readiness gates will be evaluated for pod readiness. A pod is ready when all its containers are ready AND all conditions specified in the readiness gates have status equal to "True" More info: https://git.k8s.io/enhancements/keps/sig-network/0007-pod-ready%2B%2B.md

参数	参数类型	描述
restartPolicy	String	Restart policy for all containers within the pod. Only use Always when creating a deployment. Default to Always. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle/#restart-policy
runtimeClassName	String	RuntimeClassName refers to a RuntimeClass object in the node.k8s.io group, which should be used to run this pod. If no RuntimeClass resource matches the named class, the pod will not be run. If unset or empty, the "legacy" RuntimeClass will be used, which is an implicit class with an empty definition that uses the default runtime handler. More info: https://git.k8s.io/enhancements/keps/sig-node/runtime-class.md This is a beta feature as of Kubernetes v1.14.
schedulerName	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler.
securityContext	io.k8s.api.core.v1.PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Optional: Defaults to empty. See type description for default values of each field.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
serviceAccountName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod. More info: https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/
setHostnameAsFQDN	Boolean	If true the pod's hostname will be configured as the pod's FQDN, rather than the leaf name (the default). In Linux containers, this means setting the FQDN in the hostname field of the kernel (the nodename field of struct utsname). In Windows containers, this means setting the registry value of hostname for the registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters to FQDN. If a pod does not have FQDN, this has no effect. Default to false.

参数	参数类型	描述
shareProcessNamespace	Boolean	Share a single process namespace between all of the containers in a pod. When this is set containers will be able to view and signal processes from other containers in the same pod, and the first process in each container will not be assigned PID 1. HostPID and ShareProcessNamespace cannot both be set. Optional: Default to false.
subdomain	String	If specified, the fully qualified Pod hostname will be "...svc.". If not specified, the pod will not have a domainname at all.
terminationGracePeriodSeconds	Long	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.
tolerations	Array of io.k8s.api.core.v1.Toleration objects	If specified, the pod's tolerations.
topologySpreadConstraints	Array of io.k8s.api.core.v1.TopologySpreadConstraint objects	TopologySpreadConstraints describes how a group of pods ought to spread across topology domains. Scheduler will schedule pods in a way which abides by the constraints. All topologySpreadConstraints are ANDed.
volumes	Array of io.k8s.api.core.v1.Volume objects	List of volumes that can be mounted by containers belonging to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes

表 5-6916 io.k8s.api.core.v1.Affinity

参数	参数类型	描述
nodeAffinity	io.k8s.api.core.v1.NodeAffinity object	Describes node affinity scheduling rules for the pod.

参数	参数类型	描述
podAffinity	io.k8s.api.core.v1.PodAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	io.k8s.api.core.v1.PodAntiAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 5-6917 io.k8s.api.core.v1.NodeAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PreferredSchedulingTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	io.k8s.api.core.v1.NodeSelector object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 5-6918 io.k8s.api.core.v1.PreferredSchedulingTerm

参数	参数类型	描述
preference	io.k8s.api.core.v1.NodeSelectorTerm object	A node selector term, associated with the corresponding weight.

参数	参数类型	描述
weight	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 5-6919 io.k8s.api.core.v1.NodeSelectorTerm

参数	参数类型	描述
matchExpressions	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's labels.
matchFields	Array of io.k8s.api.core.v1.NodeSelectorRequirement objects	A list of node selector requirements by node's fields.

表 5-6920 io.k8s.api.core.v1.NodeSelector

参数	参数类型	描述
nodeSelectorTerms	Array of io.k8s.api.core.v1.NodeSelectorTerm objects	Required. A list of node selector terms. The terms are ORed.

表 5-6921 io.k8s.api.core.v1.NodeSelectorRequirement

参数	参数类型	描述
key	String	The label key that the selector applies to.
operator	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist, Gt, and Lt.

参数	参数类型	描述
values	Array of strings	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 5-6922 io.k8s.api.core.v1.PodAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6923 io.k8s.api.core.v1.PodAntiAffinity

参数	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.WeightedPodAffinityTerm objects	The scheduler will prefer to schedule pods to nodes that satisfy the anti-affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling anti-affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which matches the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	Array of io.k8s.api.core.v1.PodAffinityTerm objects	If the anti-affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the anti-affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.

表 5-6924 io.k8s.api.core.v1.WeightedPodAffinityTerm

参数	参数类型	描述
podAffinityTerm	io.k8s.api.core.v1.PodAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 5-6925 io.k8s.api.core.v1.PodAffinityTerm

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over a set of resources, in this case pods.
namespaces	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"
topologyKey	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. Empty topologyKey is not allowed.

表 5-6926 io.k8s.api.core.v1.PodDNSConfig

参数	参数类型	描述
nameservers	Array of strings	A list of DNS name server IP addresses. This will be appended to the base nameservers generated from DNSPolicy. Duplicated nameservers will be removed.
options	Array of io.k8s.api.core.v1.PodDNSConfigOption objects	A list of DNS resolver options. This will be merged with the base options generated from DNSPolicy. Duplicated entries will be removed. Resolution options given in Options will override those that appear in the base DNSPolicy.
searches	Array of strings	A list of DNS search domains for host-name lookup. This will be appended to the base search paths generated from DNSPolicy. Duplicated search paths will be removed.

表 5-6927 io.k8s.api.core.v1.PodDNSConfigOption

参数	参数类型	描述
name	String	Required.
value	String	value is the value of the option

表 5-6928 io.k8s.api.core.v1.EphemeralContainer

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Lifecycle is not allowed for ephemeral containers.
livenessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
name	String	Name of the ephemeral container specified as a DNS_LABEL. This name must be unique among all containers, init containers and ephemeral containers.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	Ports are not allowed for ephemeral containers.
readinessProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources are not allowed for ephemeral containers. Ephemeral containers use spare resources already allocated to the pod.
securityContext	io.k8s.api.core.v1.SecurityContext object	SecurityContext is not allowed for ephemeral containers.
startupProbe	io.k8s.api.core.v1.Probe object	Probes are not allowed for ephemeral containers.
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	参数类型	描述
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
targetContainerName	String	If set, the name of the container from PodSpec that this ephemeral container targets. The ephemeral container will be run in the namespaces (IPC, PID, etc) of this container. If not set then the ephemeral container is run in whatever namespaces are shared for the pod. Note that the container runtime must support this feature.
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

参数	参数类型	描述
volumeDevices	Array of io.k8s.api.core.v1.Volume Device objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.Volume Mount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6929 io.k8s.api.core.v1.HostAlias

参数	参数类型	描述
hostnames	Array of strings	Hostnames for the above IP address.
ip	String	IP address of the host file entry.

表 5-6930 io.k8s.api.core.v1.LocalObjectReference

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

表 5-6931 io.k8s.api.core.v1.Container

参数	参数类型	描述
args	Array of strings	Arguments to the entrypoint. The docker image's CMD is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
command	Array of strings	Entrypoint array. Not executed within a shell. The docker image's ENTRYPOINT is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated. More info: https://kubernetes.io/docs/tasks/inject-data-application/define-command-argument-container/#running-a-command-in-a-shell
env	Array of io.k8s.api.core.v1.EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	Array of io.k8s.api.core.v1.EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.

参数	参数类型	描述
image	String	Docker image name. More info: https://kubernetes.io/docs/concepts/containers/images This field is optional to allow higher level config management to default or override container images in workload controllers like Deployments and StatefulSets.
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. More info: https://kubernetes.io/docs/concepts/containers/images#updating-images
lifecycle	io.k8s.api.core.v1.Lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.
livenessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
ports	Array of io.k8s.api.core.v1.ContainerPort objects	List of ports to expose from the container. Exposing a port here gives the system additional information about the network connections a container uses, but is primarily informational. Not specifying a port here DOES NOT prevent that port from being exposed. Any port which is listening on the default "0.0.0.0" address inside a container will be accessible from the network. Cannot be updated.
readinessProbe	io.k8s.api.core.v1.Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
resources	io.k8s.api.core.v1.ResourceRequirements object	Compute Resources required by this container. Cannot be updated. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

参数	参数类型	描述
securityContext	io.k8s.api.core.v1.SecurityContext object	Security options the pod should run with. More info: https://kubernetes.io/docs/concepts/policy/security-context/ More info: https://kubernetes.io/docs/tasks/configure-pod-container/security-context/
startupProbe	io.k8s.api.core.v1.Probe object	StartupProbe indicates that the Pod has successfully initialized. If specified, no other probes are executed until this completes successfully. If this probe fails, the Pod will be restarted, just as if the livenessProbe failed. This can be used to provide different probe parameters at the beginning of a Pod's lifecycle, when it might take a long time to load data or warm a cache, than during steady-state operation. This cannot be updated. This is a beta feature enabled by the StartupProbe feature flag. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container processes that reads from stdin will never receive an EOF. Default is false
terminationMessagePath	String	Optional: Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Will be truncated by the node if greater than 4096 bytes. The total message length across all containers will be limited to 12kb. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
terminationMessagePolicy	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.
volumeDevices	Array of io.k8s.api.core.v1.VolumeDevice objects	volumeDevices is the list of block devices to be used by the container.
volumeMounts	Array of io.k8s.api.core.v1.VolumeMount objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
workingDir	String	Container's working directory. If not specified, the container runtime's default will be used, which might be configured in the container image. Cannot be updated.

表 5-6932 io.k8s.api.core.v1.EnvVar

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, ie: \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	参数类型	描述
valueFrom	io.k8s.api.core.v1.EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 5-6933 io.k8s.api.core.v1.EnvVarSource

参数	参数类型	描述
configMapKeyRef	io.k8s.api.core.v1.ConfigMapKeySelector object	Selects a key of a ConfigMap.
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Selects a field of the pod: supports <code>metadata.name</code> , <code>metadata.namespace</code> , <code>metadata.labels['<KEY>']</code> , <code>metadata.annotations['<KEY>']</code> , <code>spec.nodeName</code> , <code>spec.serviceAccountName</code> , <code>status.hostIP</code> , <code>status.podIP</code> , <code>status.podIPs</code> .
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, limits.ephemeral-storage, requests.cpu, requests.memory and requests.ephemeral-storage) are currently supported.
secretKeyRef	io.k8s.api.core.v1.SecretKeySelector object	Selects a key of a secret in the pod's namespace

表 5-6934 io.k8s.api.core.v1.ConfigMapKeySelector

参数	参数类型	描述
key	String	The key to select.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its key must be defined

表 5-6935 io.k8s.api.core.v1.SecretKeySelector

参数	参数类型	描述
key	String	The key of the secret to select from. Must be a valid secret key.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-6936 io.k8s.api.core.v1.EnvFromSource

参数	参数类型	描述
configMapRef	io.k8s.api.core.v1.ConfigMapEnvSource object	The ConfigMap to select from
prefix	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	io.k8s.api.core.v1.SecretEnvSource object	The Secret to select from

表 5-6937 io.k8s.api.core.v1.ConfigMapEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap must be defined

表 5-6938 io.k8s.api.core.v1.SecretEnvSource

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names

参数	参数类型	描述
optional	Boolean	Specify whether the Secret must be defined

表 5-6939 io.k8s.api.core.v1.Lifecycle

参数	参数类型	描述
postStart	io.k8s.api.core.v1.Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks
preStop	io.k8s.api.core.v1.Handler object	PreStop is called immediately before a container is terminated due to an API request or management event such as liveness/startup probe failure, preemption, resource contention, etc. The handler is not called if the container crashes or exits. The reason for termination is passed to the handler. The Pod's termination grace period countdown begins before the PreStop hooked is executed. Regardless of the outcome of the handler, the container will eventually terminate within the Pod's termination grace period. Other management of the container blocks until the hook completes or until the termination grace period is reached. More info: https://kubernetes.io/docs/concepts/containers/container-lifecycle-hooks/#container-hooks

表 5-6940 io.k8s.api.core.v1.Handler

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported

表 5-6941 io.k8s.api.core.v1.ContainerPort

参数	参数类型	描述
containerPort	Integer	Number of port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
hostIP	String	What host IP to bind the external port to.
hostPort	Integer	Number of port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
protocol	String	Protocol for port. Must be UDP, TCP, or SCTP. Defaults to "TCP".

表 5-6942 io.k8s.api.core.v1.SecurityContext

参数	参数类型	描述
allowPrivilegeEscalation	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN
capabilities	io.k8s.api.core.v1.Capabilities object	The capabilities to add/drop when running containers. Defaults to the default set of capabilities granted by the container runtime.
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Defaults to false.

参数	参数类型	描述
procMount	String	procMount denotes the type of proc mount to use for the containers. The default is DefaultProcMount which uses the container runtime defaults for readonly paths and masked paths. This requires the ProcMountType feature flag to be enabled.
readOnlyRootFilesystem	Boolean	Whether this container has a read-only root filesystem. Default is false.
runAsGroup	Long	The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	Long	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by this container. If seccomp options are provided at both the pod & container level, the container options override the pod options.

参数	参数类型	描述
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options from the PodSecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6943 io.k8s.api.core.v1.Capabilities

参数	参数类型	描述
add	Array of strings	Added capabilities
drop	Array of strings	Removed capabilities

表 5-6944 io.k8s.api.core.v1.Probe

参数	参数类型	描述
exec	io.k8s.api.core.v1.ExecAction object	One and only one of the following should be specified. Exec specifies the action to take.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.
httpGet	io.k8s.api.core.v1.HTTPGetAction object	HTTPGet specifies the http request to perform.
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness and startup. Minimum value is 1.

参数	参数类型	描述
tcpSocket	io.k8s.api.core.v1.TCPSocketAction object	TCPSocket specifies an action involving a TCP port. TCP hooks not yet supported
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1. More info: https://kubernetes.io/docs/concepts/workloads/pods/pod-lifecycle#container-probes

表 5-6945 io.k8s.api.core.v1.ExecAction

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply exec'd, it is not run inside a shell, so traditional shell instructions (' ', etc) won't work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 5-6946 io.k8s.api.core.v1.HTTPGetAction

参数	参数类型	描述
host	String	Host name to connect to, defaults to the pod IP. You probably want to set "Host" in httpHeaders instead.
httpHeaders	Array of io.k8s.api.core.v1.HTTPHeader objects	Custom headers to set in the request. HTTP allows repeated headers.
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

参数	参数类型	描述
policy	io.k8s.api.core.v1.Policy object	VolumeMount Policy.
readOnly	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Defaults to false.
subPath	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).
subPathExpr	String	Expanded path within the volume from which the container's volume should be mounted. Behaves similarly to SubPath but environment variable references \$(VAR_NAME) are expanded using the container's environment. Defaults to "" (volume's root). SubPathExpr and SubPath are mutually exclusive.

表 5-6951 io.k8s.api.core.v1.Policy

参数	参数类型	描述
logs	io.k8s.api.core.v1.Logs object	Logs describes log Volume and its rotate strategy.

表 5-6952 io.k8s.api.core.v1.Logs

参数	参数类型	描述
annotations	Map<String,String>	Annotations for log.
rotate	String	Rotate strategy, including 'Daily' 'Hourly' 'Weekly'.

表 5-6953 io.k8s.api.core.v1.PodReadinessGate

参数	参数类型	描述
conditionType	String	ConditionType refers to a condition in the pod's condition list with matching type.

表 5-6954 io.k8s.api.core.v1.PodSecurityContext

参数	参数类型	描述
fsGroup	Long	<p>A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod:</p> <ol style="list-style-type: none">1. The owning GID will be the FSGroup2. The setgid bit is set (new files created in the volume will be owned by FSGroup)3. The permission bits are OR'd with rw-rw---- <p>If unset, the Kubelet will not modify the ownership and permissions of any volume.</p>
fsGroupChangePolicy	String	<p>fsGroupChangePolicy defines behavior of changing ownership and permission of the volume before being exposed inside Pod. This field will only apply to volume types which support fsGroup based ownership (and permissions). It will have no effect on ephemeral volume types such as: secret, configmaps and emptydir. Valid values are "OnRootMismatch" and "Always". If not specified defaults to "Always".</p>
runAsGroup	Long	<p>The GID to run the entrypoint of the container process. Uses runtime default if unset. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>
runAsNonRoot	Boolean	<p>Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.</p>
runAsUser	Long	<p>The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.</p>

参数	参数类型	描述
seLinuxOptions	io.k8s.api.core.v1.SELinuxOptions object	The SELinux context to be applied to all containers. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
seccompProfile	io.k8s.api.core.v1.SeccompProfile object	The seccomp options to use by the containers in this pod.
supplementalGroups	Array of longs	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
sysctls	Array of io.k8s.api.core.v1.Sysctl objects	Sysctls hold a list of namespaced sysctls used for the pod. Pods with unsupported sysctls (by the container runtime) might fail to launch.
windowsOptions	io.k8s.api.core.v1.WindowsSecurityContextOptions object	The Windows specific settings applied to all containers. If unspecified, the options within a container's SecurityContext will be used. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6955 io.k8s.api.core.v1.SELinuxOptions

参数	参数类型	描述
level	String	Level is SELinux level label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
user	String	User is a SELinux user label that applies to the container.

表 5-6956 io.k8s.api.core.v1.SeccompProfile

参数	参数类型	描述
localhostProfile	String	localhostProfile indicates a profile defined in a file on the node should be used. The profile must be preconfigured on the node to work. Must be a descending path, relative to the kubelet's configured seccomp profile location. Must only be set if type is "Localhost".
type	String	type indicates which kind of seccomp profile will be applied. Valid options are: Localhost - a profile defined in a file on the node should be used. RuntimeDefault - the container runtime default profile should be used. Unconfined - no profile should be applied.

表 5-6957 io.k8s.api.core.v1.Sysctl

参数	参数类型	描述
name	String	Name of a property to set
value	String	Value of a property to set

表 5-6958 io.k8s.api.core.v1.WindowsSecurityContextOptions

参数	参数类型	描述
gmsaCredentialSpec	String	GMSACredentialSpec is where the GMSA admission webhook (https://github.com/kubernetes-sigs/windows-gmsa) inlines the contents of the GMSA credential spec named by the GMSACredentialSpecName field.
gmsaCredentialSpecName	String	GMSACredentialSpecName is the name of the GMSA credential spec to use.
runAsUserName	String	The UserName in Windows to run the entrypoint of the container process. Defaults to the user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 5-6959 io.k8s.api.core.v1.Toleration

参数	参数类型	描述
effect	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	Long	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 5-6960 io.k8s.api.core.v1.TopologySpreadConstraint

参数	参数类型	描述
labelSelector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	LabelSelector is used to find matching pods. Pods that match this label selector are counted to determine the number of pods in their corresponding topology domain.

参数	参数类型	描述
maxSkew	Integer	MaxSkew describes the degree to which pods may be unevenly distributed. When <i>whenUnsatisfiable=DoNotSchedule</i> , it is the maximum permitted difference between the number of matching pods in the target topology and the global minimum. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 1/1/0: zone1 zone2 zone3 P P - if MaxSkew is 1, incoming pod can only be scheduled to zone3 to become 1/1/1; scheduling it onto zone1(zone2) would make the ActualSkew(2-0) on zone1(zone2) violate MaxSkew(1). - if MaxSkew is 2, incoming pod can be scheduled onto any zone. When <i>whenUnsatisfiable=ScheduleAnyway</i> , it is used to give higher precedence to topologies that satisfy it. It's a required field. Default value is 1 and 0 is not allowed.
topologyKey	String	TopologyKey is the key of node labels. Nodes that have a label with this key and identical values are considered to be in the same topology. We consider each <key, value> as a "bucket", and try to put balanced number of pods into each bucket. It's a required field.
whenUnsatisfiable	String	WhenUnsatisfiable indicates how to deal with a pod if it doesn't satisfy the spread constraint. - DoNotSchedule (default) tells the scheduler not to schedule it. - ScheduleAnyway tells the scheduler to schedule the pod in any location, but giving higher precedence to topologies that would help reduce the skew. A constraint is considered "Unsatisfiable" for an incoming pod if and only if every possible node assignment for that pod would violate "MaxSkew" on some topology. For example, in a 3-zone cluster, MaxSkew is set to 1, and pods with the same labelSelector spread as 3/1/1: zone1 zone2 zone3 P P P P P If WhenUnsatisfiable is set to DoNotSchedule, incoming pod can only be scheduled to zone2(zone3) to become 3/2/1(3/1/2) as ActualSkew(2-1) on zone2(zone3) satisfies MaxSkew(1). In other words, the cluster can still be imbalanced, but scheduler won't make it <i>more</i> imbalanced. It's a required field.

表 5-6961 io.k8s.api.core.v1.Volume

参数	参数类型	描述
awsElasticBlockStore	io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource object	AWSElasticBlockStore represents an AWS Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
azureDisk	io.k8s.api.core.v1.AzureDiskVolumeSource object	AzureDisk represents an Azure Data Disk mount on the host and bind mount to the pod.
azureFile	io.k8s.api.core.v1.AzureFileVolumeSource object	AzureFile represents an Azure File Service mount on the host and bind mount to the pod.
cephfs	io.k8s.api.core.v1.CephFSVolumeSource object	CephFS represents a Ceph FS mount on the host that shares a pod's lifetime
cinder	io.k8s.api.core.v1.CinderVolumeSource object	Cinder represents a cinder volume attached and mounted on kubelets host machine. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
configMap	io.k8s.api.core.v1.ConfigMapVolumeSource object	ConfigMap represents a configMap that should populate this volume
csi	io.k8s.api.core.v1.CSIVolumeSource object	CSI (Container Storage Interface) represents ephemeral storage that is handled by certain external CSI drivers (Beta feature).
downwardAPI	io.k8s.api.core.v1.DownwardAPIVolumeSource object	DownwardAPI represents downward API about the pod that should populate this volume
emptyDir	io.k8s.api.core.v1.EmptyDirVolumeSource object	EmptyDir represents a temporary directory that shares a pod's lifetime. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir

参数	参数类型	描述
ephemeral	io.k8s.api.core.v1.EphemeralVolumeSource object	<p>Ephemeral represents a volume that is handled by a cluster storage driver (Alpha feature). The volume's lifecycle is tied to the pod that defines it - it will be created before the pod starts, and deleted when the pod is removed.</p> <p>Use this if: a) the volume is only needed while the pod runs, b) features of normal volumes like restoring from snapshot or capacity tracking are needed, c) the storage driver is specified through a storage class, and d) the storage driver supports dynamic volume provisioning through a PersistentVolumeClaim (see EphemeralVolumeSource for more information on the connection between this volume type and PersistentVolumeClaim).</p> <p>Use PersistentVolumeClaim or one of the vendor-specific APIs for volumes that persist for longer than the lifecycle of an individual pod.</p> <p>Use CSI for light-weight local ephemeral volumes if the CSI driver is meant to be used that way - see the documentation of the driver for more information.</p> <p>A pod can use both types of ephemeral volumes and persistent volumes at the same time.</p>
fc	io.k8s.api.core.v1.FCVolumeSource object	FC represents a Fibre Channel resource that is attached to a kubelet's host machine and then exposed to the pod.
flexVolume	io.k8s.api.core.v1.FlexVolumeSource object	FlexVolume represents a generic volume resource that is provisioned/attached using an exec based plugin.
flocker	io.k8s.api.core.v1.FlockerVolumeSource object	Flocker represents a Flocker volume attached to a kubelet's host machine. This depends on the Flocker control service being running
gcePersistentDisk	io.k8s.api.core.v1.GCEPersistentDiskVolumeSource object	GCEPersistentDisk represents a GCE Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
gitRepo	io.k8s.api.core.v1.GitRepoVolumeSource object	GitRepo represents a git repository at a particular revision. DEPRECATED: GitRepo is deprecated. To provision a container with a git repo, mount an EmptyDir into an InitContainer that clones the repo using git, then mount the EmptyDir into the Pod's container.
glusterfs	io.k8s.api.core.v1.GlusterfsVolumeSource object	Glusterfs represents a Glusterfs mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/glusterfs/README.md
hostPath	io.k8s.api.core.v1.HostPathVolumeSource object	HostPath represents a pre-existing file or directory on the host machine that is directly exposed to the container. This is generally used for system agents or other privileged things that are allowed to see the host machine. Most containers will NOT need this. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
iscsi	io.k8s.api.core.v1.ISCSIVolumeSource object	ISCSI represents an ISCSI Disk resource that is attached to a kubelet's host machine and then exposed to the pod. More info: https://examples.k8s.io/volumes/iscsi/README.md
localDir	io.k8s.api.core.v1.LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
name	String	Volume's name. Must be a DNS_LABEL and unique within the pod. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
nfs	io.k8s.api.core.v1.NFSVolumeSource object	NFS represents an NFS mount on the host that shares a pod's lifetime More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
persistentVolumeClaim	io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
photonPersistentDisk	io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource object	PhotonPersistentDisk represents a PhotonController persistent disk attached and mounted on kubelets host machine

参数	参数类型	描述
portworxVolume	io.k8s.api.core.v1.PortworxVolumeSource object	PortworxVolume represents a portworx volume attached and mounted on kubelets host machine
projected	io.k8s.api.core.v1.ProjectedVolumeSource object	Items for all in one resources secrets, configmaps, and downward API
quobyte	io.k8s.api.core.v1.QuobyteVolumeSource object	Quobyte represents a Quobyte mount on the host that shares a pod's lifetime
rbd	io.k8s.api.core.v1.RBDVolumeSource object	RBD represents a Rados Block Device mount on the host that shares a pod's lifetime. More info: https://examples.k8s.io/volumes/rbd/README.md
scaleIO	io.k8s.api.core.v1.ScaleIOVolumeSource object	ScaleIO represents a ScaleIO persistent volume attached and mounted on Kubernetes nodes.
secret	io.k8s.api.core.v1.SecretVolumeSource object	Secret represents a secret that should populate this volume. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret
storageos	io.k8s.api.core.v1.StorageOSVolumeSource object	StorageOS represents a StorageOS volume attached and mounted on Kubernetes nodes.
vsphereVolume	io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource object	VsphereVolume represents a vSphere volume attached and mounted on kubelets host machine

表 5-6962 io.k8s.api.core.v1.AWSElasticBlockStoreVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty).
readOnly	Boolean	Specify "true" to force and set the ReadOnly property in VolumeMounts to "true". If omitted, the default is "false". More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore
volumeID	String	Unique ID of the persistent disk resource in AWS (Amazon EBS volume). More info: https://kubernetes.io/docs/concepts/storage/volumes#awselasticblockstore

表 5-6963 io.k8s.api.core.v1.AzureDiskVolumeSource

参数	参数类型	描述
cachingMode	String	Host Caching mode: None, Read Only, Read Write.
diskName	String	The Name of the data disk in the blob storage
diskURI	String	The URI the data disk in the blob storage
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
kind	String	Expected values Shared: multiple blob disks per storage account Dedicated: single blob disk per storage account Managed: azure managed data disk (only in managed availability set). defaults to shared

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.

表 5-6964 io.k8s.api.core.v1.AzureFileVolumeSource

参数	参数类型	描述
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretName	String	the name of secret that contains Azure Storage Account Name and Key
shareName	String	Share Name

表 5-6965 io.k8s.api.core.v1.CephFSVolumeSource

参数	参数类型	描述
monitors	Array of strings	Required: Monitors is a collection of Ceph monitors More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
path	String	Optional: Used as the mounted root, rather than the full Ceph tree, default is /
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretFile	String	Optional: SecretFile is the path to key ring for User, default is /etc/ceph/user.secret More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the authentication secret for User, default is empty. More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it
user	String	Optional: User is the rados user name, default is admin More info: https://examples.k8s.io/volumes/cephfs/README.md#how-to-use-it

表 5-6966 io.k8s.api.core.v1.CinderVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts. More info: https://examples.k8s.io/mysql-cinder-pd/README.md
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: points to a secret object containing parameters used to connect to OpenStack.
volumeID	String	volume id used to identify the volume in cinder. More info: https://examples.k8s.io/mysql-cinder-pd/README.md

表 5-6967 io.k8s.api.core.v1.ConfigMapVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.

参数	参数类型	描述
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6968 io.k8s.api.core.v1.CSIVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the CSI driver that handles this volume. Consult with your admin for the correct name as registered in the cluster.
fsType	String	Filesystem type to mount. Ex. "ext4", "xfs", "ntfs". If not provided, the empty value is passed to the associated CSI driver which will determine the default filesystem to apply.
nodePublishSecretRef	io.k8s.api.core.v1.LocalObjectReference object	NodePublishSecretRef is a reference to the secret object containing sensitive information to pass to the CSI driver to complete the CSI NodePublishVolume and NodeUnpublishVolume calls. This field is optional, and may be empty if no secret is required. If the secret object contains more than one secret, all secret references are passed.
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).
volumeAttributes	Map<String,String>	VolumeAttributes stores driver-specific properties that are passed to the CSI driver. Consult your driver's documentation for supported values.

表 5-6969 io.k8s.api.core.v1.DownwardAPIVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits to use on created files by default. Must be a Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of downward API volume file

表 5-6970 io.k8s.api.core.v1.EmptyDirVolumeSource

参数	参数类型	描述
medium	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory. More info: https://kubernetes.io/docs/concepts/storage/volumes#emptydir
sizeLimit	String	Total amount of local storage required for this EmptyDir volume. The size limit is also applicable for memory medium. The maximum usage on memory medium EmptyDir would be the minimum value between the SizeLimit specified here and the sum of memory limits of all containers in a pod. The default is nil which means that the limit is undefined. More info: http://kubernetes.io/docs/user-guide/volumes#emptydir

表 5-6971 io.k8s.api.core.v1.EphemeralVolumeSource

参数	参数类型	描述
readOnly	Boolean	Specifies a read-only configuration for the volume. Defaults to false (read/write).

参数	参数类型	描述
volumeClaimTemplate	io.k8s.api.core.v1.PersistentVolumeClaimTemplate object	<p>Will be used to create a stand-alone PVC to provision the volume. The pod in which this EphemeralVolumeSource is embedded will be the owner of the PVC, i.e. the PVC will be deleted together with the pod. The name of the PVC will be <i><pod name>-<volume name></i> where <i><volume name></i> is the name from the <i>PodSpec.Volumes</i> array entry. Pod validation will reject the pod if the concatenated name is not valid for a PVC (for example, too long).</p> <p>An existing PVC with that name that is not owned by the pod will <i>not</i> be used for the pod to avoid using an unrelated volume by mistake. Starting the pod is then blocked until the unrelated PVC is removed. If such a pre-created PVC is meant to be used by the pod, the PVC has to updated with an owner reference to the pod once the pod exists. Normally this should not be necessary, but it may be useful when manually reconstructing a broken cluster.</p> <p>This field is read-only and no changes will be made by Kubernetes to the PVC after it has been created.</p> <p>Required, must not be nil.</p>

表 5-6972 io.k8s.api.core.v1.PersistentVolumeClaimTemplate

参数	参数类型	描述
metadata	io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta object	May contain labels and annotations that will be copied into the PVC when creating it. No other fields are allowed and will be rejected during validation.
spec	io.k8s.api.core.v1.PersistentVolumeClaimSpec object	The specification for the PersistentVolumeClaim. The entire content is copied unchanged into the PVC that gets created from this template. The same fields as in a PersistentVolumeClaim are also valid here.

表 5-6973 io.k8s.apimachinery.pkg.apis.meta.v1.ObjectMeta

参数	参数类型	描述
annotations	Map<String,String>	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. More info: http://kubernetes.io/docs/user-guide/annotations
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata
deletionGracePeriodSeconds	Long	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
deletionTimestamp	String	<p>DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource is expected to be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field, once the finalizers list is empty. As long as the finalizers list contains items, deletion is blocked. Once the deletionTimestamp is set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. After that 30 seconds, the Kubelet will send a hard termination signal (SIGKILL) to the container and after cleanup, remove the pod from the API. In the presence of network partitions, this object may still exist after this timestamp, until an administrator or automated process can determine the resource is fully terminated. If not set, graceful deletion of the object has not been requested.</p> <p>Populated by the system when a graceful deletion is requested. Read-only. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#metadata</p>

参数	参数类型	描述
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed. Finalizers may be processed and removed in any order. Order is NOT enforced because it introduces significant risk of stuck finalizers. finalizers is a shared field, any actor with permission can reorder it. If the finalizer list is processed in order, then this can lead to a situation in which the component responsible for the first finalizer in the list is waiting for a signal (field value, external system, or other) produced by a component responsible for a finalizer later in the list, resulting in a deadlock. Without enforced ordering finalizers are free to order amongst themselves and are not vulnerable to ordering changes in the list.
generateName	String	GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server. If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header). Applied only if Name is not specified. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#idempotency
generation	Long	A sequence number representing a specific generation of the desired state. Populated by the system. Read-only.

参数	参数类型	描述
labels	Map<String,String>	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services. More info: http://kubernetes.io/docs/user-guide/labels
managedFields	Array of io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry objects	ManagedFields maps workflow-id and version to the set of fields that are managed by that workflow. This is mostly for internal housekeeping, and users typically shouldn't need to set or understand this field. A workflow can be the user's name, a controller's name, or the name of a specific apply path like "ci-cd". The set of fields is always in the version that the workflow used when modifying the object.
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/identifiers#names
namespace	String	Namespace defines the space within which each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. More info: http://kubernetes.io/docs/user-guide/namespaces
ownerReferences	Array of io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.

参数	参数类型	描述
resourceVersion	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources.</p> <p>Populated by the system. Read-only. Value must be treated as opaque by clients and . More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#concurrency-control-and-consistency</p>
selfLink	String	<p>SelfLink is a URL representing this object. Populated by the system. Read-only.</p> <p>DEPRECATED Kubernetes will stop propagating this field in 1.20 release and the field is planned to be removed in 1.21 release.</p>
uid	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations.</p> <p>Populated by the system. Read-only. More info: http://kubernetes.io/docs/user-guide/identifiers#uids</p>

表 5-6974 io.k8s.apimachinery.pkg.apis.meta.v1.ManagedFieldsEntry

参数	参数类型	描述
apiVersion	String	<p>APIVersion defines the version of this resource that this field set applies to. The format is "group/version" just like the top-level APIVersion field. It is necessary to track the version of a field set because it cannot be automatically converted.</p>
fieldsType	String	<p>FieldsType is the discriminator for the different fields format and version. There is currently only one possible value: "FieldsV1"</p>
fieldsV1	Object	<p>FieldsV1 holds the first JSON version format as described in the "FieldsV1" type.</p>

参数	参数类型	描述
manager	String	Manager is an identifier of the workflow managing these fields.
operation	String	Operation is the type of operation which lead to this ManagedFieldsEntry being created. The only valid values for this field are 'Apply' and 'Update'.
time	String	Time is timestamp of when these fields were set. It should always be empty if Operation is 'Apply'

表 5-6975 io.k8s.apimachinery.pkg.apis.meta.v1.OwnerReference

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
controller	Boolean	If true, this reference points to the managing controller.
kind	String	Kind of the referent. More info: https://git.k8s.io/community/contributors/devel/sig-architecture/api-conventions.md#types-kinds
name	String	Name of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#names
uid	String	UID of the referent. More info: http://kubernetes.io/docs/user-guide/identifiers#uids

表 5-6976 io.k8s.api.core.v1.PersistentVolumeClaimSpec

参数	参数类型	描述
accessModes	Array of strings	AccessModes contains the desired access modes the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#access-modes-1

参数	参数类型	描述
dataSource	io.k8s.api.core.v1.TypedLocalObjectReference object	This field can be used to specify either: * An existing VolumeSnapshot object (snapshot.storage.k8s.io/VolumeSnapshot - Beta) * An existing PVC (PersistentVolumeClaim) * An existing custom resource/object that implements data population (Alpha) In order to use VolumeSnapshot object types, the appropriate feature gate must be enabled (VolumeSnapshotDataSource or AnyVolumeDataSource) If the provisioner or an external controller can support the specified data source, it will create a new volume based on the contents of the specified data source. If the specified data source is not supported, the volume will not be created and the failure will be reported as an event. In the future, we plan to support more data source types and the behavior of the provisioner may change.
resources	io.k8s.api.core.v1.ResourceRequirements object	Resources represents the minimum resources the volume should have. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#resources
selector	io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector object	A label query over volumes to consider for binding.
storageClassName	String	Name of the StorageClass required by the claim. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#class-1
volumeMode	String	volumeMode defines what type of volume is required by the claim. Value of Filesystem is implied when not included in claim spec.
volumeName	String	VolumeName is the binding reference to the PersistentVolume backing this claim.

表 5-6977 io.k8s.api.core.v1.TypedLocalObjectReference

参数	参数类型	描述
apiGroup	String	APIGroup is the group for the resource being referenced. If APIGroup is not specified, the specified Kind must be in the core API group. For any other third-party types, APIGroup is required.
kind	String	Kind is the type of resource being referenced
name	String	Name is the name of resource being referenced

表 5-6978 io.k8s.api.core.v1.ResourceRequirements

参数	参数类型	描述
limits	Map<String,String>	Limits describes the maximum amount of compute resources allowed. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/
requests	Map<String,String>	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. More info: https://kubernetes.io/docs/concepts/configuration/manage-compute-resources-container/

表 5-6979 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelector

参数	参数类型	描述
matchExpressions	Array of io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	Map<String,String>	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 5-6980 io.k8s.apimachinery.pkg.apis.meta.v1.LabelSelectorRequirement

参数	参数类型	描述
key	String	key is the label key that the selector applies to.
operator	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 5-6981 io.k8s.api.core.v1.FCVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
lun	Integer	Optional: FC target lun number
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
targetWWNs	Array of strings	Optional: FC target worldwide names (WWNs)
wwids	Array of strings	Optional: FC volume world wide identifiers (wwids) Either wwids or combination of targetWWNs and lun must be set, but not both simultaneously.

表 5-6982 io.k8s.api.core.v1.FlexVolumeSource

参数	参数类型	描述
driver	String	Driver is the name of the driver to use for this volume.
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". The default filesystem depends on FlexVolume script.

参数	参数类型	描述
options	Map<String,String>	Optional: Extra command options if any.
readOnly	Boolean	Optional: Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	Optional: SecretRef is reference to the secret object containing sensitive information to pass to the plugin scripts. This may be empty if no secret object is specified. If the secret object contains more than one secret, all secrets are passed to the plugin scripts.

表 5-6983 io.k8s.api.core.v1.FlockerVolumeSource

参数	参数类型	描述
datasetName	String	Name of the dataset stored as metadata -> name on the dataset for Flocker should be considered as deprecated
datasetUUID	String	UUID of the dataset. This is unique identifier of a Flocker dataset

表 5-6984 io.k8s.api.core.v1.GCEPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
partition	Integer	The partition in the volume that you want to mount. If omitted, the default is to mount by volume name. Examples: For volume /dev/sda1, you specify the partition as "1". Similarly, the volume partition for /dev/sda is "0" (or you can leave the property empty). More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

参数	参数类型	描述
pdName	String	Unique name of the PD resource in GCE. Used to identify the disk in GCE. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#gcepersistentdisk

表 5-6985 io.k8s.api.core.v1.GitRepoVolumeSource

参数	参数类型	描述
directory	String	Target directory name. Must not contain or start with '..'. If '.' is supplied, the volume directory will be the git repository. Otherwise, if specified, the volume will contain the git repository in the subdirectory with the given name.
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 5-6986 io.k8s.api.core.v1.GlusterfsVolumeSource

参数	参数类型	描述
endpoints	String	EndpointsName is the endpoint name that details Glusterfs topology. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
path	String	Path is the Glusterfs volume path. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod
readOnly	Boolean	ReadOnly here will force the Glusterfs volume to be mounted with read-only permissions. Defaults to false. More info: https://examples.k8s.io/volumes/glusterfs/README.md#create-a-pod

表 5-6987 io.k8s.api.core.v1.HostPathVolumeSource

参数	参数类型	描述
path	String	Path of the directory on the host. If the path is a symlink, it will follow the link to the real path. More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath
type	String	Type for HostPath Volume Defaults to "" More info: https://kubernetes.io/docs/concepts/storage/volumes#hostpath

表 5-6988 io.k8s.api.core.v1.ISCSIVolumeSource

参数	参数类型	描述
chapAuthDiscovery	Boolean	whether support iSCSI Discovery CHAP authentication
chapAuthSession	Boolean	whether support iSCSI Session CHAP authentication
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#iscsi
initiatorName	String	Custom iSCSI Initiator Name. If initiatorName is specified with iscsiInterface simultaneously, new iSCSI interface : will be created for the connection.
iqn	String	Target iSCSI Qualified Name.
iscsiInterface	String	iSCSI Interface Name that uses an iSCSI transport. Defaults to 'default' (tcp).
lun	Integer	iSCSI Target Lun number.
portals	Array of strings	iSCSI Target Portal List. The portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	CHAP Secret for iSCSI target and initiator authentication

参数	参数类型	描述
targetPortal	String	iSCSI Target Portal. The Portal is either an IP or ip_addr:port if the port is other than default (typically TCP ports 860 and 3260).

表 5-6989 io.k8s.api.core.v1.LocalDirVolumeSource

参数	参数类型	描述
sizeLimit	String	<p>Quantity is a fixed-point representation of a number. It provides convenient marshaling/unmarshaling in JSON and YAML, in addition to String() and AsInt64() accessors.</p> <p>The serialization format is:</p> <p>::= (Note that may be empty, from the "" case in .) ::= 0 1 ... 9 ::= ::= . . . ::= "+" "-" ::= ::= ::= Ki Mi Gi Ti Pi Ei (International System of units; See: http://physics.nist.gov/cuu/Units/binary.html) ::= m "" k M G T P E (Note that 1024 = 1Ki but 1000 = 1k; I didn't choose the capitalization.) ::= "e" "E"</p> <p>No matter which of the three exponent forms is used, no quantity may represent a number greater than 2⁶³-1 in magnitude, nor may it have more than 3 decimal places. Numbers larger or more precise will be capped or rounded up. (E.g.: 0.1m will rounded up to 1m.) This may be extended in the future if we require larger or smaller quantities.</p> <p>When a Quantity is parsed from a string, it will remember the type of suffix it had, and will use the same type again when it is serialized.</p> <p>Before serializing, Quantity will be put in "canonical form". This means that Exponent/suffix will be adjusted up or down (with a corresponding increase or decrease in Mantissa) such that: a. No precision is lost b. No fractional digits will be emitted c. The exponent (or suffix) is as large as possible. The sign will be omitted unless the number is negative.</p> <p>Examples: 1.5 will be serialized as "1500m" 1.5Gi will be serialized as "1536Mi"</p> <p>Note that the quantity will NEVER be internally represented by a floating point number. That is the whole point of this exercise.</p> <p>Non-canonical values will still parse as long as they are well formed, but will be re-emitted in their canonical form. (So always use canonical form, or don't diff.)</p> <p>This format is intended to make it difficult to use these numbers without writing some sort of special handling code in the hopes that that</p>

参数	参数类型	描述
		will cause implementors to also use a fixed point implementation.

表 5-6990 io.k8s.api.core.v1.NFSVolumeSource

参数	参数类型	描述
path	String	Path that is exported by the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
readOnly	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs
server	String	Server is the hostname or IP address of the NFS server. More info: https://kubernetes.io/docs/concepts/storage/volumes#nfs

表 5-6991 io.k8s.api.core.v1.PersistentVolumeClaimVolumeSource

参数	参数类型	描述
claimName	String	ClaimName is the name of a PersistentVolumeClaim in the same namespace as the pod using this volume. More info: https://kubernetes.io/docs/concepts/storage/persistent-volumes#persistentvolumeclaims
readOnly	Boolean	Will force the ReadOnly setting in VolumeMounts. Default false.

表 5-6992 io.k8s.api.core.v1.PhotonPersistentDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
pdID	String	ID that identifies Photon Controller persistent disk

表 5-6993 io.k8s.api.core.v1.PortworxVolumeSource

参数	参数类型	描述
fsType	String	FSType represents the filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
volumeID	String	VolumeID uniquely identifies a Portworx volume

表 5-6994 io.k8s.api.core.v1.ProjectedVolumeSource

参数	参数类型	描述
defaultMode	Integer	Mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
sources	Array of io.k8s.api.core.v1.VolumeProjection objects	list of volume projections

表 5-6995 io.k8s.api.core.v1.VolumeProjection

参数	参数类型	描述
configMap	io.k8s.api.core.v1.ConfigMapProjection object	information about the configMap data to project
downwardAPI	io.k8s.api.core.v1.DownwardAPIProjection object	information about the downwardAPI data to project

参数	参数类型	描述
secret	io.k8s.api.core.v1.SecretProjection object	information about the secret data to project
serviceAccountToken	io.k8s.api.core.v1.ServiceAccountTokenProjection object	information about the serviceAccountToken data to project

表 5-6996 io.k8s.api.core.v1.ConfigMapProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced ConfigMap will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the ConfigMap, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the ConfigMap or its keys must be defined

表 5-6997 io.k8s.api.core.v1.DownwardAPIProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.DownwardAPIVolumeFile objects	Items is a list of DownwardAPIVolume file

表 5-6998 io.k8s.api.core.v1.DownwardAPIVolumeFile

参数	参数类型	描述
fieldRef	io.k8s.api.core.v1.ObjectFieldSelector object	Required: Selects a field of the pod: only annotations, labels, name and namespace are supported.
mode	Integer	Optional: mode bits used to set permissions on this file, must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
resourceFieldRef	io.k8s.api.core.v1.ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 5-6999 io.k8s.api.core.v1.ObjectFieldSelector

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1".
fieldPath	String	Path of the field to select in the specified API version.

表 5-7000 io.k8s.api.core.v1.ResourceFieldSelector

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars
divisor	String	Specifies the output format of the exposed resources, defaults to "1"
resource	String	Required: resource to select

表 5-7001 io.k8s.api.core.v1.SecretProjection

参数	参数类型	描述
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
name	String	Name of the referent. More info: https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#names
optional	Boolean	Specify whether the Secret or its key must be defined

表 5-7002 io.k8s.api.core.v1.ServiceAccountTokenProjection

参数	参数类型	描述
audience	String	Audience is the intended audience of the token. A recipient of a token must identify itself with an identifier specified in the audience of the token, and otherwise should reject the token. The audience defaults to the identifier of the apiserver.
expirationSeconds	Long	ExpirationSeconds is the requested duration of validity of the service account token. As the token approaches expiration, the kubelet volume plugin will proactively rotate the service account token. The kubelet will start trying to rotate the token if the token is older than 80 percent of its time to live or if the token is older than 24 hours. Defaults to 1 hour and must be at least 10 minutes.
path	String	Path is the path relative to the mount point of the file to project the token into.

表 5-7003 io.k8s.api.core.v1.QuobyteVolumeSource

参数	参数类型	描述
group	String	Group to map volume access to Default is no group
readOnly	Boolean	ReadOnly here will force the Quobyte volume to be mounted with read-only permissions. Defaults to false.
registry	String	Registry represents a single or multiple Quobyte Registry services specified as a string as host:port pair (multiple entries are separated with commas) which acts as the central registry for volumes
tenant	String	Tenant owning the given Quobyte volume in the Backend Used with dynamically provisioned Quobyte volumes, value is set by the plugin
user	String	User to map volume access to Defaults to serviceaccount user
volume	String	Volume is a string that references an already created Quobyte volume by name.

表 5-7004 io.k8s.api.core.v1.RBDVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type of the volume that you want to mount. Tip: Ensure that the filesystem type is supported by the host operating system. Examples: "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified. More info: https://kubernetes.io/docs/concepts/storage/volumes#rbd
image	String	The rados image name. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
keyring	String	Keyring is the path to key ring for RBDUser. Default is /etc/ceph/keyring. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
monitors	Array of strings	A collection of Ceph monitors. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

参数	参数类型	描述
pool	String	The rados pool name. Default is rbd. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
readOnly	Boolean	ReadOnly here will force the ReadOnly setting in VolumeMounts. Defaults to false. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef is name of the authentication secret for RBDUser. If provided overrides keyring. Default is nil. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it
user	String	The rados user name. Default is admin. More info: https://examples.k8s.io/volumes/rbd/README.md#how-to-use-it

表 5-7005 io.k8s.api.core.v1.ScaleIOVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Default is "xfs".
gateway	String	The host address of the ScaleIO API Gateway.
protectionDomain	String	The name of the ScaleIO Protection Domain for the configured storage.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef references to the secret for ScaleIO user and other sensitive information. If this is not provided, Login operation will fail.
sslEnabled	Boolean	Flag to enable/disable SSL communication with Gateway, default false
storageMode	String	Indicates whether the storage for a volume should be ThickProvisioned or ThinProvisioned. Default is ThinProvisioned.
storagePool	String	The ScaleIO Storage Pool associated with the protection domain.

参数	参数类型	描述
system	String	The name of the storage system as configured in ScaleIO.
volumeName	String	The name of a volume already created in the ScaleIO system that is associated with this volume source.

表 5-7006 io.k8s.api.core.v1.SecretVolumeSource

参数	参数类型	描述
defaultMode	Integer	Optional: mode bits used to set permissions on created files by default. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
items	Array of io.k8s.api.core.v1.KeyToPath objects	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error unless it is marked optional. Paths must be relative and may not contain the '..' path or start with '..'.
optional	Boolean	Specify whether the Secret or its keys must be defined
secretName	String	Name of the secret in the pod's namespace to use. More info: https://kubernetes.io/docs/concepts/storage/volumes#secret

表 5-7007 io.k8s.api.core.v1.KeyToPath

参数	参数类型	描述
key	String	The key to project.

参数	参数类型	描述
mode	Integer	Optional: mode bits used to set permissions on this file. Must be an octal value between 0000 and 0777 or a decimal value between 0 and 511. YAML accepts both octal and decimal values, JSON requires decimal values for mode bits. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
path	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '..'. May not start with the string '..'.

表 5-7008 io.k8s.api.core.v1.StorageOSVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
readOnly	Boolean	Defaults to false (read/write). ReadOnly here will force the ReadOnly setting in VolumeMounts.
secretRef	io.k8s.api.core.v1.LocalObjectReference object	SecretRef specifies the secret to use for obtaining the StorageOS API credentials. If not specified, default values will be attempted.
volumeName	String	VolumeName is the human-readable name of the StorageOS volume. Volume names are only unique within a namespace.
volumeNamespace	String	VolumeNamespace specifies the scope of the volume within StorageOS. If no namespace is specified then the Pod's namespace will be used. This allows the Kubernetes name scoping to be mirrored within StorageOS for tighter integration. Set VolumeName to any name to override the default behaviour. Set to "default" if you are not using namespaces within StorageOS. Namespaces that do not pre-exist within StorageOS will be created.

表 5-7009 io.k8s.api.core.v1.VsphereVirtualDiskVolumeSource

参数	参数类型	描述
fsType	String	Filesystem type to mount. Must be a filesystem type supported by the host operating system. Ex. "ext4", "xfs", "ntfs". Implicitly inferred to be "ext4" if unspecified.
storagePolicyID	String	Storage Policy Based Management (SPBM) profile ID associated with the StoragePolicyName.
storagePolicyName	String	Storage Policy Based Management (SPBM) profile name.
volumePath	String	Path that identifies vSphere volume vmdk

表 5-7010 io.k8s.api.apps.v1.ReplicaSetStatus

参数	参数类型	描述
availableReplicas	Integer	The number of available replicas (ready for at least minReadySeconds) for this replica set.
conditions	Array of io.k8s.api.apps.v1.ReplicaSetCondition objects	Represents the latest available observations of a replica set's current state.
fullyLabeledReplicas	Integer	The number of pods that have labels matching the labels of the pod template of the replicaset.
observedGeneration	Long	ObservedGeneration reflects the generation of the most recently observed ReplicaSet.
readyReplicas	Integer	The number of ready replicas for this replica set.
replicas	Integer	Replicas is the most recently observed number of replicas. More info: https://kubernetes.io/docs/concepts/workloads/controllers/replicationcontroller/#what-is-a-replicationcontroller

表 5-7011 io.k8s.api.apps.v1.ReplicaSetCondition

参数	参数类型	描述
lastTransition Time	String	The last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of replica set condition.

请求示例

无

响应示例

状态码： 200

OK

```
{
  "apiVersion": "apps/v1",
  "kind": "ReplicaSet",
  "metadata": {
    "annotations": {
      "deployment.kubernetes.io/desired-replicas": "2",
      "deployment.kubernetes.io/max-replicas": "2",
      "deployment.kubernetes.io/revision": "1"
    },
    "creationTimestamp": "2022-09-06T07:42:15Z",
    "generation": 1,
    "labels": {
      "app": "deployment-test",
      "pod-template-hash": "86cc979d6d"
    },
    "name": "replicaset-test",
    "namespace": "namespace-test",
    "ownerReferences": [ {
      "apiVersion": "apps/v1",
      "blockOwnerDeletion": true,
      "controller": true,
      "kind": "Deployment",
      "name": "deployment-test",
      "uid": "1e5c7f22-36b3-4a1d-868e-5fc10a98f332"
    } ],
    "resourceVersion": "41587974",
    "selfLink": "/apis/apps/v1/namespaces/namespace-test/replicasets/replicaset-test",
    "uid": "03037364-044e-43ca-bc68-06d07ab0d609"
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "matchLabels": {
        "app": "deployment-test",
        "pod-template-hash": "86cc979d6d"
      }
    }
  }
}
```

```
}
},
"template": {
  "metadata": {
    "annotations": {
      "cri.cci.io/container-type": "secure-container",
      "log.stdoutcollection.kubernetes.io": "{\"collectionContainers\": [\"container-0\"]}",
      "metrics.alpha.kubernetes.io/custom-endpoints": "[{api: \"\", path: \"\", port: \"\", names: \"\"}]"
    },
    "creationTimestamp": null,
    "labels": {
      "app": "deployment-test",
      "pod-template-hash": "86cc979d6d"
    }
  },
  "spec": {
    "containers": [ {
      "image": "redis",
      "imagePullPolicy": "IfNotPresent",
      "lifecycle": { },
      "name": "container-0",
      "resources": {
        "limits": {
          "cpu": "500m",
          "memory": "1Gi"
        },
        "requests": {
          "cpu": "500m",
          "memory": "1Gi"
        }
      },
      "terminationMessagePath": "/dev/termination-log",
      "terminationMessagePolicy": "File"
    } ],
    "dnsConfig": { },
    "dnsPolicy": "ClusterFirst",
    "imagePullSecrets": [ {
      "name": "imagepull-secret"
    } ],
    "restartPolicy": "Always",
    "schedulerName": "default-scheduler",
    "securityContext": { },
    "terminationGracePeriodSeconds": 30
  }
}
},
"status": {
  "availableReplicas": 2,
  "fullyLabeledReplicas": 2,
  "observedGeneration": 1,
  "readyReplicas": 2,
  "replicas": 2
}
}
```

状态码

状态码	描述
200	OK
400	BadRequest
401	Unauthorized
403	Forbidden

状态码	描述
404	NotFound
405	MethodNotAllowed
406	NotAcceptable
409	Conflict
415	UnsupportedMediaType
422	Invalid
429	TooManyRequests
500	InternalServerError
503	ServiceUnavailable
504	ServerTimeout

6 数据结构

本章描述API使用的公共参数。

6.1 请求数据结构（废弃）

表 6-1 v1.PodTemplate 数据结构说明

参数	是否必选	参数类型	描述
kind	Yes	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is PodTemplate .
apiVersion	Yes	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	Yes	metadata object	-
template	Yes	template object	-

表 6-2 v1.Pod 请求参数

参数	是否必选	参数类型	描述
kind	Yes	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Pod .
apiVersion	Yes	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	Yes	metadata object	-
spec	Yes	spec object	-
status	No	status object	-

表 6-3 status 字段数据结构说明

参数	是否必选	参数类型	描述
phase	No	String	Current condition of the pod.
conditions	No	conditions object	Current service state of the pod.
message	No	String	A human readable message indicating details about why the pod is in this condition.
reason	No	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'OutOfDisk'
hostIP	No	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.

参数	是否必选	参数类型	描述
podIP	No	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
startTime	No	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.
containerStatuses	No	containerStatuses object	The list has one entry per container in the manifest. Each entry is currently the output of container inspect.

表 6-4 conditions 字段数据结构说明

参数	是否必选	参数类型	描述
type	No	String	Type of the condition. Currently only Ready.
status	No	String	Status of the condition. Can be True, False, or Unknown.
lastProbeTime	No	String	Last time we probed the condition.
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.
reason	No	String	Unique, one-word, CamelCase reason for the condition's last transition.
message	No	String	Human-readable message indicating details about last transition.

表 6-5 containerStatuses 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
state	No	state/lastState object	-
lastState	No	state/lastState object	-
ready	No	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	No	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. However, those containers are subject to garbage collection. This value will get capped at 5 by GC.
image	Yes	String	The image the container is running.
imageID	No	String	ID of the container's image.
containerID	No	String	Container's ID in the format 'docker://'.

表 6-6 state/lastState 字段数据结构说明

参数	是否必选	参数类型	描述
waiting	No	waiting object	-
running	No	running object	-
terminated	No	terminated object	-

表 6-7 waiting 字段数据结构说明

参数	是否必选	参数类型	描述
reason	No	String	(Brief) Reason the container is not yet running.
message	No	String	Message regarding why the container is not yet running.

表 6-8 running 字段数据结构说明

参数	是否必选	参数类型	描述
startedAt	No	String	Time at which the container was last (re-)started.

表 6-9 terminated 字段数据结构说明

参数	是否必选	参数类型	描述
startedAt	No	Integer	Exit status from the last termination of the container.
signal	No	Integer	Signal from the last termination of the container.
reason	No	String	(Brief) reason from the last termination of the container.
message	No	String	Message regarding the last termination of the container.
startedAt	No	String	Time at which previous execution of the container started.
finishedAt	No	String	Time at which the container last terminated.
containerID	No	String	Container's ID in the format 'docker://'

表 6-10 metadata 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	<p>Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated.</p> <p>0 characters < name length ≤ 253 characters.</p> <p>The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.</p>
clusterName	No	String	<p>The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.</p> <p>说明 除了evs卷，其他类型的卷不支持</p>

参数	是否必选	参数类型	描述
initializers	No	initializers object	An initializer is a controller which enforces some system invariant at object creation time. This field is a list of initializers that have not yet acted on this object. If nil or empty, this object has been completely initialized. Otherwise, the object is considered uninitialized and is hidden (in list/watch and get calls) from clients that haven't explicitly asked to observe uninitialized objects. When an object is created, the system will populate this list with the current set of initializers. Only privileged users may set or modify this list. Once it is empty, it may not be modified further by any user.

参数	是否必选	参数类型	描述
generateName	No	String	<p>An optional prefix used by the server to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different from the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409. Instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified.</p> <p>0 characters < generated name length ≤ 253 characters.</p> <p>The generated name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.</p>

参数	是否必选	参数类型	描述
namespace	No	String	Namespace defines the space within each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated. 0 characters < namespace length ≤ 63 characters. The namespace must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
selfLink	No	String	A URL representing this object. Populated by the system. Read-only. 说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.
uid	No	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only. 说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.

参数	是否必选	参数类型	描述
resourceVersion	No	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
generation	No	Integer	<p>A sequence number representing a specific generation of the desired state. Currently only implemented by replication controllers. Populated by the system. Read-only.</p>
creationTimestamp	No	String	<p>A timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>

参数	是否必选	参数类型	描述
deletionTimestamp	No	String	RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource will be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field. Once set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. Once the resource is deleted in the API, the Kubelet will send a hard termination signal to the container. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only.
deletionGracePeriodSeconds	No	Integer	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
labels	Yes	Map[string]string	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services.

参数	是否必选	参数类型	描述
Annotations	No	Map[string]string	An unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects.
ownerReferences	No	ownerReferences object	(A newly added parameter in Kubernetes 1.3) List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
finalizers	No	Array of strings	(A newly added parameter in Kubernetes 1.3) Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed.

表 6-11 initializers 字段数据结构说明

参数	是否必选	参数类型	描述
pending	No	pending object	Pending is a list of initializers that must execute in order before this object is visible. When the last pending initializer is removed, and no failing result is set, the initializers struct will be set to nil and the object is considered as initialized and visible to all clients.
result	No	result object	If result is set with the Failure field, the object will be persisted to storage and then deleted, ensuring that other clients can observe the deletion.

表 6-12 pending 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	name of the process that is responsible for initializing this object.

表 6-13 result 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
code	No	Integer	Suggested HTTP return code for this status, 0 if not set.

参数	是否必选	参数类型	描述
details	No	details object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
message	No	String	A human-readable description of the status of this operation.
metadata	Yes	metadata object	Standard list metadata.
reason	No	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	No	String	Status of the operation. One of: "Success" or "Failure".

表 6-14 details 字段数据结构说明

参数	是否必选	参数类型	描述
causes	No	causes object	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	No	String	The group attribute of the resource associated with the status StatusReason.

参数	是否必选	参数类型	描述
kind	No	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind.
name	No	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	No	Integer	If specified, the time in seconds before the operation should be retried.
uid	No	String	UID of the resource. (when there is a single resource which can be described).

表 6-15 metadata 字段数据结构说明

参数	是否必选	参数类型	描述
resourceVersion	No	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only.
selfLink	No	String	SelfLink is a URL representing this object. Populated by the system. Read-only.

表 6-16 causes 字段数据结构说明

参数	是否必选	参数类型	描述
field	No	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	No	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	No	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 6-17 ownerReferences 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	API version of the referent.
blockOwnerDeletion	No	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Default to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
kind	Yes	String	Kind of the referent.

参数	是否必选	参数类型	描述
name	Yes	String	Name of the referent.
uid	No	String	UID of the referent.
controller	No	Boolean	If true, this reference points to the managing controller.

表 6-18 spec 字段数据结构说明

参数	是否必选	参数类型	描述
replicas	No	Integer	The number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Value range: ≥ 0 . Default: 1
minReadySeconds	No	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its containers crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
template	Yes	template object	-
selector	Yes	Map[string]string	A label query over pods that should match the Replicas count. If Selector is empty, it is defaulted to the labels present on the Pod template. Label keys and values that must match in order to be controlled by this replication controller, if empty defaulted to labels on Pod template.

表 6-19 status 字段数据结构说明

参数	是否必选	参数类型	描述
replicas	No	Integer	The most recently observed number of replicas.

参数	是否必选	参数类型	描述
availableReplicas	No	Integer	The number of available replicas (ready for at least minReadySeconds) for this replication controller.
readyReplicas	No	Integer	The number of ready replicas for this replication controller.
conditions	No	condition object	Represents the latest available observations of a replication controller's current state.
observedGeneration	No	Integer	Reflects the generation of the most recently observed replication controller.
FullylabeledReplicas	No	Map[string]string	-

表 6-20 template 字段数据结构说明

参数	是否必选	参数类型	描述
metadata	No	metadata object	-
spec	Yes	spec object	-

表 6-21 condition 字段数据结构说明

参数	是否必选	参数类型	描述
lastTransitionTime	No	Time	The last time the condition transitioned from one status to another.
message	No	String	A human readable message indicating details about the transition.
reason	No	String	The reason for the condition's last transition.
status	No	String	Status of the condition, one of True, False, Unknown.
type	No	String	Type of replication controller condition.

表 6-22 spec 字段数据结构说明

参数	是否必选	参数类型	描述
volumes	No	-	Not supported now.
affinity	No	affinity object	If specified, the pod's scheduling constraints 说明 不允许用户设置affinity，默认使用软反亲和
containers	Yes	containers object	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a pod. Cannot be updated.
restartPolicy	No	String	Restart policy for all containers within the pod. Value: <ul style="list-style-type: none">• Always• OnFailure• Never Default: Always.
activeDeadlineSeconds	No	Integer	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. Value range of this parameter: > 0.
dnsPolicy	No	String	Set DNS policy for containers within the pod. Value: <ul style="list-style-type: none">• ClusterFirst• Default Default: ClusterFirst.
hostAliases	No	hostAliases object	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.

参数	是否必选	参数类型	描述
serviceAccount Name	No	String	Name of the ServiceAccount used to run this pod. 0 characters < service account name length ≤ 253 characters. The service account name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?. 说明 不支持serviceaccount,所以不支持该字段
serviceAccount	No	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead.
schedulerName	No	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. 说明 不支持设置指定scheduler name
nodeName	No	String	A request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements. 0 characters < node name length ≤ 253 characters. The node name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?. 说明 不支持设置nodeName
nodeSelector	No	Object	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. 说明 不支持设置nodeSelector

参数	是否必选	参数类型	描述
automountServiceAccountToken	No	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
hostNetwork	No	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. 说明 The hostport of the hostNetwork must be different from the containerport. 创建、更新pod时不支持使用主机网络
hostPID	No	Boolean	A flag indicating whether to use the host's pid namespace. Optional: Default to false. 说明 The hostport of the hostNetwork must be different from the containerport. 创建、更新pod时不支持使用主机PID namespaces
hostIPC	No	Boolean	A flag indicating whether to use the host's ipc namespace. Optional: Default to false. 说明 The hostport of the hostNetwork must be different from the containerport. 创建、更新pod时不支持使用主机IPC namespaces
securityContext	No	securityContext object	-
imagePullSecrets	No	imagePullSecrets object	A list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. 说明 使用的镜像是容器镜像服务界面上的“我的镜像”页签中的镜像时必须设置。

参数	是否必选	参数类型	描述
initContainers	No	containers object	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, or Liveness probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed.
hostname	No	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
subdomain	No	String	If specified, the fully qualified Pod hostname will be "<hostname>.<subdomain>.<pod namespace>.svc<cluster domain>". If not specified, the pod will not have a domainname at all.
tolerations	No	tolerations object	If specified, the pod's tolerations.

表 6-23 containers 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. 0 characters < container name length ≤ 63 characters. The container name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
image	Yes	String	Container image name.
command	No	Array of strings	Entrypoint array. Not executed within a shell. The container image's entrypoint is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated.

参数	是否必选	参数类型	描述
args	No	Array of strings	Arguments to the entrypoint. The container image's cmd is used if this is not provided. Variable references \$ (VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$ (VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated.
workingDir	No	String	Container's working directory. Defaults to Container's default. Defaults to image's default. Cannot be updated.
ports	No	ports object	List of ports to expose from the container. Cannot be updated.
env	No	env object	List of environment variables to set in the container. Cannot be updated.
envFrom	No	envFrom object	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
resources	No	resources object	Minimum resources the volume should have.

参数	是否必选	参数类型	描述
volumeMounts	No	volumeMounts object	Pod volumes to mount into the container's filesystem. Cannot be updated.
livenessProbe	No	livenessProbe object	-
readinessProbe	No	livenessProbe object	-
lifecycle	No	lifecycle object	-
terminationMessagePath	No	String	Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Defaults to /dev/termination-log. Cannot be updated.
imagePullPolicy	No	String	Image pull policy. Defaults to Always if the :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. Value: <ul style="list-style-type: none">• Always• Never• IfNotPresent
securityContext	No	securityContext object	-
stdin	No	Boolean	A flag indicating whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.

参数	是否必选	参数类型	描述
stdinOnce	No	Boolean	A flag indicating whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true, the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container process that reads from stdin will never receive an EOF. Default is false.
tty	No	Boolean	A flag indicating whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

表 6-24 securityContext 字段数据结构说明

参数	是否必选	参数类型	描述
seLinuxOptions	No	seLinuxOptions object	-
runAsUser	No	Integer	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Value length: > 0 characters.

参数	是否必选	参数类型	描述
runAsNonRoot	No	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
supplementalGroups	No	Array of integers	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
fsGroup	No	Integer	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw.

表 6-25 imagePullSecrets 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the referent. 须知 创建应用时，使用的镜像是容器镜像服务界面上的“我的镜像”页签中的镜像，那么该参数的值必须设置为imagepull-secret。

表 6-26 hostPath 字段数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Path of the directory on the host.

表 6-27 emptyDir 字段数据结构说明

参数	是否必选	参数类型	描述
medium	No	String	What type of storage medium should back this directory. The default is "" which means to use the node's default medium. Must be an empty string (default) or Memory.

表 6-28 gitRepo 字段数据结构说明

参数	是否必选	参数类型	描述
repository	No	String	Repository URL.
revision	No	String	Commit hash for the specified revision.

表 6-29 ports 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. 0 characters < name length ≤ 15 characters. The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.

参数	是否必选	参数类型	描述
hostPort	No	Integer	Number of the port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. Value range: [1, 65535].
containerPort	No	Integer	Number of the port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$. Value range: [1, 65535].
protocol	No	String	Protocol for port. Value: <ul style="list-style-type: none">• TCP• UDP Default: TCP.
hostIP	No	String	What host IP to bind the external port to.

表 6-30 env 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	No	String	Variable references $\$(VAR_NAME)$ are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The $\$(VAR_NAME)$ syntax can be escaped with a double $\$\$$, for example, $\$\(VAR_NAME) . Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".

参数	是否必选	参数类型	描述
valueFrom	No	valueFrom object	-

表 6-31 resources 字段数据结构说明

参数	是否必选	参数类型	描述
limits	No	Object	Maximum amount of compute resources allowed.
requests	No	Object	Minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. 云容器实例中Pod规格有限制，具体的限制请参见 约束限制 页面的“Pod规格”部分。

表 6-32 volumeMounts 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	This must match the Name of a Volume. 0 character < name length ≤ 253 characters. The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
readOnly	No	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Value: <ul style="list-style-type: none">• true• false Default: false.
mountPath	No	String	Path within the container at which the volume should be mounted. Value length: > 0 characters.

参数	是否必选	参数类型	描述
subPath	No	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

表 6-33 livenessProbe 字段数据结构说明

参数	是否必选	参数类型	描述
exec	No	exec object	-
httpGet	No	httpGet object	-
tcpSocket	No	tcpSocket object	-
initialDelaySeconds	No	Integer	Number of seconds after the container has started before liveness probes are initiated. Value range: ≥ 0 .
timeoutSeconds	No	Integer	Number of seconds after which the probe times out. Value range: ≥ 0 . Default: 1.
periodSeconds	No	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. Value range: ≥ 0 . Default: 10.
successThreshold	No	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness. Minimum value is 1. Value range: ≥ 0 . Default: 1.

参数	是否必选	参数类型	描述
failureThreshold	No	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. Value range: ≥ 0 . Default: 3.

表 6-34 lifecycle 字段数据结构说明

参数	是否必选	参数类型	描述
postStart	No	postStart/ preStop object	-
preStop	No	postStart/ preStop object	-

表 6-35 securityContext 字段数据结构说明

参数	是否必选	参数类型	描述
capabilities	No	capabilities object	-
privileged	No	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Value: <ul style="list-style-type: none">• true• false Default: false. 说明 Run container in privileged mode
seLinuxOptions	No	seLinuxOptions object	-

参数	是否必选	参数类型	描述
runAsUser	No	Integer	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	No	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Value: <ul style="list-style-type: none">• true• false

表 6-36 seLinuxOptions 字段数据结构说明

参数	是否必选	参数类型	描述
user	No	String	SELinux user label that applies to the container.
role	No	String	SELinux role label that applies to the container.
type	No	String	SELinux type label that applies to the container.
level	No	String	SELinux level label that applies to the container.

表 6-37 items 字段数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
fieldRef	No	fieldRef object	-
resourceFieldRef	No	resourceFieldRef object	Selects a resource of the container: only resources limits and requests. (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 6-38 valueFrom 字段数据结构说明

参数	是否必选	参数类型	描述
fieldRef	No	fieldRef object	-
resourceFieldRef	No	resourceFieldRef object	Selects a resource of the container: only resources limits and requests. (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 6-39 exec 字段数据结构说明

参数	是否必选	参数类型	描述
command	No	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply executed, it is not run inside a shell, so traditional shell instructions (' ', etc) do not work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 6-40 httpGet 字段数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Path to access on the HTTP server.
port	Yes	String	Name or number of the port to access on the container. The port number must be in the range 1 to 65535. The port name must be an IANA_SVC_NAME.
host	No	String	Host name to connect to. Defaults to the pod IP address.
scheme	No	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 6-41 tcpSocket 字段数据结构说明

参数	是否必选	参数类型	描述
port	Yes	String	Number or name of the port to access on the container. The port number must be in the range 1 to 65535. The port name must be an IANA_SVC_NAME.

表 6-42 postStart/preStop 字段数据结构说明

参数	是否必选	参数类型	描述
exec	No	exec object	-
httpGet	No	httpGet object	-
tcpSocket	No	tcpSocket object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported.

表 6-43 capabilities 字段数据结构说明

参数	是否必选	参数类型	描述
add	No	add object	Added capabilities.
drop	No	add object	Removed capabilities.

表 6-44 fieldRef 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	No	String	Version of the schema the FieldPath is written in terms of. Defaults to "v1".
fieldPath	No	String	Path of the field to select in the specified API version.

表 6-45 resourceFieldRef 字段数据结构说明

参数	是否必选	参数类型	描述
containerName	No	String	Container name: required for volumes, optional for env vars.
resource	Yes	String	Required: resource to select.
divisor	No	Integer	Specifies the output format of the exposed resources, defaults to "1".

表 6-46 add 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the resource.
namespaced	No	Boolean	A flag indicating whether a resource is namespaced or not. Default: false.
kind	No	String	kind of the resource.

表 6-47 affinity 字段数据结构说明

参数	是否必选	参数类型	描述
nodeAffinity	No	nodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	No	podAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	No	podAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 6-48 nodeAffinity 字段数据结构说明

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	No	preferredDuringSchedulingIgnoredDuringExecution object	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.
requiredDuringSchedulingIgnoredDuringExecution	No	requiredDuringSchedulingIgnoredDuringExecution object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 6-49 podAffinity 字段数据结构说明

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	No	preferredDuringSchedulingIgnoredDuringExecution object	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which match the corresponding podAffinityTerm; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	No	podAffinityTerm object	<p>NOT YET IMPLEMENTED. TODO: Uncomment field once it is implemented. If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system will try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.</p> <p><code>RequiredDuringSchedulingRequiredDuringExecution []PodAffinityTerm</code> json:"requiredDuringSchedulingRequiredDuringExecution,omitempty" If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each <code>podAffinityTerm</code> are intersected, i.e. all terms must be satisfied.</p>

表 6-50 preferredDuringSchedulingIgnoredDuringExecution 字段数据结构说明

参数	是否必选	参数类型	描述
preference	No	preference object	A node selector term, associated with the corresponding weight.
weight	No	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 6-51 requiredDuringSchedulingIgnoredDuringExecution 字段数据结构说明

参数	是否必选	参数类型	描述
nodeSelectorTerms	No	preference object	Required. A list of node selector terms. The terms are ORed.

表 6-52 preference 字段数据结构说明

参数	是否必选	参数类型	描述
matchExpressions	No	matchExpressions object	Required. A list of node selector requirements. The requirements are ANDed.

表 6-53 matchExpressions 字段数据结构说明

参数	是否必选	参数类型	描述
key	No	String	The label key that the selector applies to.
operator	No	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.

参数	是否必选	参数类型	描述
values	No	String	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 6-54 preferredDuringSchedulingIgnoredDuringExecution 字段数据结构说明

参数	是否必选	参数类型	描述
podAffinityTerm	No	podAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	No	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 6-55 podAffinityTerm 字段数据结构说明

参数	是否必选	参数类型	描述
labelSelector	No	labelSelector object	A label query over a set of resources, in this case pods.
namespaces	No	Array[string]	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"

参数	是否必选	参数类型	描述
topologyKey	No	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. For PreferredDuringScheduling pod anti-affinity, empty topologyKey is interpreted as "all topologies" ("all topologies" here means all the topologyKeys indicated by scheduler command-line argument --failure-domains); for affinity and for RequiredDuringScheduling pod anti-affinity, empty topologyKey is not allowed.

表 6-56 labelSelector 字段数据结构说明

参数	是否必选	参数类型	描述
matchExpressions	No	matchExpressions object	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	No	Object	(A newly added parameter in Kubernetes 1.7) matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 6-57 matchExpressions 字段数据结构说明

参数	是否必选	参数类型	描述
key	No	String	key is the label key that the selector applies to.
operator	No	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	No	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 6-58 hostAliases 字段数据结构说明

参数	是否必选	参数类型	描述
hostnames	No	Array of strings	Hostnames for the above IP address.
ip	No	String	IP address of the host file entry.

表 6-59 envFrom 字段数据结构说明

参数	是否必选	参数类型	描述
configMapRef	No	configMapRef object	The ConfigMap to select from
prefix	No	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
secretRef	No	secretRef object	-

表 6-60 configMapRef 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	Name of the referent.
optional	No	Boolean	Specify whether the ConfigMap must be defined

表 6-61 secretRef 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	Name of the referent.
optional	No	Boolean	Specify whether the ConfigMap must be defined

表 6-62 tolerations 字段数据结构说明

参数	是否必选	参数类型	描述
effect	No	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	No	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.
operator	No	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.

参数	是否必选	参数类型	描述
tolerationSeconds	No	Integer	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	No	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 6-63 DeleteOptions 数据结构说明

参数	是否必选	参数类型	描述
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is DeleteOptions.
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1.

参数	是否必选	参数类型	描述
gracePeriodSeconds	No	Integer	The duration in seconds before the object should be deleted. Value must be a non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	No	Boolean	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	No	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy.
preconditions	No	Preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.

表 6-64 preconditions 字段数据结构说明

参数	是否必选	参数类型	描述
uid	No	String	Specifies the target UID.

6.2 响应数据结构（废弃）

表 6-65 Pod 响应参数

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	metadata object	-
spec	spec object	-
status	status object	-

表 6-66 PodList 响应参数

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	metadata object	-
items	Array of Pod objects	List of pods.

表 6-67 PodTemplate 响应参数

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	metadata object	-
template	template object	-

表 6-68 PodTemplateList 响应参数

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	metadata object	-
items	items object	List of pod templates.

表 6-69 items 字段数据结构说明

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	metadata object	-
spec	spec object	-
status	status object	-

表 6-70 status 字段数据结构说明

参数	参数类型	描述
phase	String	Current condition of the pod.
conditions	conditions object	Current service state of the pod.
message	String	A human readable message indicating details about why the pod is in this condition.
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'OutOfDisk'
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.

参数	参数类型	描述
containerStatuses	containerStatuses object	The list has one entry per container in the manifest. Each entry is currently the output of container inspect.

表 6-71 conditions 字段数据结构说明

参数	参数类型	描述
type	String	Type of the condition. Currently only Ready.
status	String	Status of the condition. Can be True, False, or Unknown.
lastProbeTime	String	Last time we probed the condition.
lastTransitionTime	String	Last time the condition transitioned from one status to another.
reason	String	Unique, one-word, CamelCase reason for the condition's last transition.
message	String	Human-readable message indicating details about last transition.

表 6-72 containerStatuses 字段数据结构说明

参数	参数类型	描述
name	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.
state	state/lastState object	-
lastState	state/lastState object	-
ready	Boolean	A flag indicating whether the container has passed its readiness probe.

参数	参数类型	描述
restartCount	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. However, those containers are subject to garbage collection. This value will get capped at 5 by GC.
image	String	Image that the container is running.
imageID	String	ID of the container's image.
containerID	String	Container's ID in the format 'docker://'.

表 6-73 state/lastState 字段数据结构说明

参数	参数类型	描述
waiting	waiting object	-
running	running object	-
terminated	terminated object	-

表 6-74 waiting 字段数据结构说明

参数	参数类型	描述
reason	String	(Brief) reason the container is not yet running.
message	String	Message regarding why the container is not yet running.

表 6-75 running 字段数据结构说明

参数	参数类型	描述
startedAt	String	Time at which the container was last (re-)started.

表 6-76 terminated 字段数据结构说明

参数	参数类型	描述
exitCode	Integer	Exit status from the last termination of the container.
signal	Integer	Signal from the last termination of the container.
reason	String	(Brief) reason from the last termination of the container.
message	String	Message regarding the last termination of the container.
startedAt	String	Time at which previous execution of the container started.
finishedAt	String	Time at which the container last terminated.
containerID	String	Container's ID in the format 'docker://'.

表 6-77 metadata 字段数据结构说明

参数	参数类型	描述
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only.

表 6-78 items 字段数据结构说明

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	metadata object	-
spec	spec object	-
status	status object	-

表 6-79 metadata 字段数据结构说明

参数	参数类型	描述
name	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated.
clusterName	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.

参数	参数类型	描述
initializers	initializers object	<p>An initializer is a controller which enforces some system invariant at object creation time. This field is a list of initializers that have not yet acted on this object. If nil or empty, this object has been completely initialized. Otherwise, the object is considered uninitialized and is hidden (in list/watch and get calls) from clients that haven't explicitly asked to observe uninitialized objects. When an object is created, the system will populate this list with the current set of initializers. Only privileged users may set or modify this list. Once it is empty, it may not be modified further by any user.</p>
generateName	String	<p>GenerateName is an optional prefix, used by the server, to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different from the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409 - instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified.</p>

参数	参数类型	描述
namespace	String	Namespace defines the space within each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated.
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.
uid	String	UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only.
resourceVersion	String	An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients.
generation	Integer	A sequence number representing a specific generation of the desired state. Currently only implemented by replication controllers. Populated by the system. Read-only.

参数	参数类型	描述
creationTimestamp	String	CreationTimestamp is a timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists.
deletionTimestamp	String	DeletionTimestamp is RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource will be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field. Once set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. Once the resource is deleted in the API, the Kubelet will send a hard termination signal to the container. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only.
deletionGracePeriodSeconds	Integer	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.

参数	参数类型	描述
labels	Map[string]string	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services.
Annotations	Map[string]string	Annotations is an unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects.
ownerReferences	ownerReferences object	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
finalizers	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed.

表 6-80 spec 字段数据结构说明

参数	参数类型	描述
replicas	Integer	Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1.

参数	参数类型	描述
selector	Object	Selector is a label query over pods that should match the Replicas count. If Selector is empty, it is defaulted to the labels present on the Pod template. Label keys and values that must match in order to be controlled by this replication controller, if empty defaulted to labels on Pod template.
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its containers crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
template	template object	-

表 6-81 status 字段数据结构说明

参数	参数类型	描述
replicas	Integer	Replicas is the most recently observed number of replicas.
availableReplicas	Integer	The number of available replicas (ready for at least minReadySeconds) for this replication controller.
readyReplicas	Integer	The number of ready replicas for this replication controller.
fullyLabeledReplicas	Integer	The number of pods that have labels matching the labels of the pod template of the replication controller.
conditions	ReplicationControllerCondition object	Represents the latest available observations of a replication controller's current state.
observedGeneration	Integer	ObservedGeneration reflects the generation of the most recently observed replication controller.

表 6-82 initializers 字段数据结构说明

参数	参数类型	描述
pending	pending object	Pending is a list of initializers that must execute in order before this object is visible. When the last pending initializer is removed, and no failing result is set, the initializers struct will be set to nil and the object is considered as initialized and visible to all clients.
result	result object	If result is set with the Failure field, the object will be persisted to storage and then deleted, ensuring that other clients can observe the deletion.

表 6-83 pending 字段数据结构说明

参数	参数类型	描述
name	String	name of the process that is responsible for initializing this object.

表 6-84 result 字段数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
code	Integer	Suggested HTTP return code for this status, 0 if not set.
details	details object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
message	String	A human-readable description of the status of this operation.
metadata	metadata object	Standard list metadata.
reason	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure".

表 6-85 details 字段数据结构说明

参数	参数类型	描述
causes	causes object	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind.
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).

参数	参数类型	描述
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried.
uid	String	UID of the resource. (when there is a single resource which can be described).

表 6-86 metadata 字段数据结构说明

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only.
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.

表 6-87 causes 字段数据结构说明

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.

参数	参数类型	描述
reason	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 6-88 ownerReferences 字段数据结构说明

参数	参数类型	描述
apiVersion	String	API version of the referent.
blockOwnerDeletion	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Default to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
kind	String	Kind of the referent.
name	String	Name of the referent.
uid	String	UID of the referent.
controller	Boolean	If true, this reference points to the managing controller.

表 6-89 template 字段数据结构说明

参数	参数类型	描述
metadata	metadata object	-
spec	spec object	-

表 6-90 spec 字段数据结构说明

参数	参数类型	描述
volumes	-	Not supported now.

参数	参数类型	描述
containers	containers object	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a Pod. Cannot be updated.
restartPolicy	String	Restart policy for all containers within the pod. One of Always, OnFailure, Never. Default to Always.
activeDeadlineSeconds	Integer	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer.
dnsPolicy	String	Set DNS policy for containers within the pod. One of 'ClusterFirst' or 'Default'. Defaults to "ClusterFirst".
serviceName	String	ServiceAccountName is the name of the ServiceAccount to use to run this pod.
serviceAccount	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceName instead.
nodeName	String	nodeName is a request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements.
hostNetwork	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false.
hostPID	Boolean	Use the host's pid namespace. Optional: Default to false.
hostIPC	Boolean	Use the host's ipc namespace. Optional: Default to false.

参数	参数类型	描述
securityContext	securityContext object	-
imagePullSecrets	imagePullSecrets object	ImagePullSecrets is an optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use.

表 6-91 containers 字段数据结构说明

参数	参数类型	描述
name	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated.
image	String	Container image name.
command	Array of strings	Entrypoint array. Not executed within a shell. The container image's entrypoint is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated.

参数	参数类型	描述
args	Array of strings	Arguments to the endpoint. The container image's cmd is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated.
workingDir	String	Container's working directory. Defaults to Container's default. Defaults to image's default. Cannot be updated.
ports	ports object	List of ports to expose from the container. Cannot be updated.
env	env object	List of environment variables to set in the container. Cannot be updated.
resources	resources object	Resources represents the minimum resources the volume should have.
volumeMounts	-	Not supported now.
livenessProbe	livenessProbe object	-
readinessProbe	livenessProbe object	-
lifecycle	lifecycle object	-
terminationMessagePath	String	Path at which the file to which the container's termination message will be written is mounted into the container's file system. Message written is intended to be brief final status, such as an assertion failure message. Defaults to /dev/termination-log. Cannot be updated.

参数	参数类型	描述
imagePullPolicy	String	Image pull policy. One of Always, Never, IfNotPresent. Defaults to Always if the latest tag is specified, or IfNotPresent otherwise. Cannot be updated.
securityContext	securityContext object	-
stdin	Boolean	Whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	Boolean	Whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true, the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container process that reads from stdin will never receive an EOF. Default is false.
tty	Boolean	Whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

表 6-92 securityContext 字段数据结构说明

参数	参数类型	描述
seLinuxOptions	seLinuxOptions object	-

参数	参数类型	描述
runAsUser	Integer	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
supplementalGroups	Array of integers	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
fsGroup	Integer	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw.

表 6-93 imagePullSecrets 字段数据结构说明

参数	参数类型	描述
name	String	Name of the referent.

表 6-94 hostPath 字段数据结构说明

参数	参数类型	描述
path	String	Path of the directory on the host.

表 6-95 gitRepo 字段数据结构说明

参数	参数类型	描述
repository	String	Repository URL
revision	String	Commit hash for the specified revision.

表 6-96 secret 字段数据结构说明

参数	参数类型	描述
secretName	String	SecretName is the name of a secret in the pod's namespace.

表 6-97 ports 字段数据结构说明

参数	参数类型	描述
name	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services.
hostPort	Integer	Number of the port to expose on the host. If specified, this must be a valid port number, $0 < x < 65536$. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this.
containerPort	Integer	Number of the port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$.
protocol	String	Protocol for port. Must be UDP or TCP. Defaults to "TCP".
hostIP	String	What host IP to bind the external port to.

表 6-98 env 字段数据结构说明

参数	参数类型	描述
name	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	String	Variable references \$(VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	valueFrom object	-

表 6-99 resources 字段数据结构说明

参数	参数类型	描述
limits	Object	Limits describes the maximum amount of compute resources allowed.
requests	Object	Requests describes the minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value.

表 6-100 livenessProbe 字段数据结构说明

参数	参数类型	描述
exec	exec object	-
httpGet	httpGet object	-
tcpSocket	tcpSocket object	-

参数	参数类型	描述
initialDelaySeconds	Integer	Number of seconds after the container has started before liveness probes are initiated.
timeoutSeconds	Integer	Number of seconds after which the probe times out. Defaults to 1 second. Minimum value is 1.
periodSeconds	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1.
successThreshold	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness. Minimum value is 1.
failureThreshold	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1.

表 6-101 lifecycle 字段数据结构说明

参数	参数类型	描述
postStart	postStart/preStop object	-
preStop	postStart/preStop object	-

表 6-102 securityContext 字段数据结构说明

参数	参数类型	描述
capabilities	capabilities object	-
privileged	Boolean	Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host. Default to false.
seLinuxOptions	seLinuxOptions object	-

参数	参数类型	描述
runAsUser	Integer	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.

表 6-103 seLinuxOptions 字段数据结构说明

参数	参数类型	描述
user	String	User is a SELinux user label that applies to the container.
role	String	Role is a SELinux role label that applies to the container.
type	String	Type is a SELinux type label that applies to the container.
level	String	Level is SELinux level label that applies to the container.

表 6-104 items 字段数据结构说明

参数	参数类型	描述
path	String	Required: Path is the relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
fieldRef	fieldRef object	-

表 6-105 valueFrom 字段数据结构说明

参数	参数类型	描述
fieldRef	fieldRef object	-
resourceFieldRef	resourceFieldRef object	Selects a resource of the container: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 6-106 exec 字段数据结构说明

参数	参数类型	描述
command	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply executed, it is not run inside a shell, so traditional shell instructions (' ', etc) do not work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 6-107 httpGet 字段数据结构说明

参数	参数类型	描述
path	String	Path to access on the HTTP server.
port	String	Name or number of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.
host	String	Host name to connect to, defaults to the pod IP.
scheme	String	Scheme to use for connecting to the host. Defaults to HTTP.

表 6-108 tcpSocket 字段数据结构说明

参数	参数类型	描述
port	String	Number or name of the port to access on the container. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME.

表 6-109 postStart/preStop 字段数据结构说明

参数	参数类型	描述
exec	exec object	-
httpGet	httpGet object	-
tcpSocket	tcpSocket object	TCP socket specifies an action involving a TCP port. TCP hooks not yet supported.

表 6-110 capabilities 字段数据结构说明

参数	参数类型	描述
add	add object	Added capabilities
drop	add object	Removed capabilities

表 6-111 fieldRef 字段数据结构说明

参数	参数类型	描述
apiVersion	String	Version of the schema the FieldPath is written in terms of, defaults to "v1"
fieldPath	String	Path of the field to select in the specified API version.

表 6-112 resourceFieldRef 字段数据结构说明

参数	参数类型	描述
containerName	String	Container name: required for volumes, optional for env vars.
resource	String	Required: resource to select.
divisor	Integer	Specifies the output format of the exposed resources, defaults to "1".

表 6-113 add 字段数据结构说明

参数	参数类型	描述
name	String	name is the name of the resource.
namespaced	Boolean	namespaced indicates if a resource is namespaced or not. Default: false.
kind	String	kind is the kind for the resource.

表 6-114 metadata 字段数据结构说明

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only.

参数	参数类型	描述
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.

表 6-115 ReplicationControllerCondition 字段数据结构说明

参数	参数类型	描述
lastTransitionTime	String	The last time the condition transitioned from one status to another.
message	String	A human readable message indicating details about the transition.
reason	String	The reason for the condition's last transition.
status	String	Status of the condition, one of True, False, Unknown.
type	String	Type of replication controller condition.

表 6-116 status v1 Core 字段数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
code	Integer	Suggested HTTP return code for this status, 0 if not set
details	StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase
message	String	A human-readable description of the status of this operation

参数	参数类型	描述
metadata	ListMeta v1 meta object	Standard list metadata.
reason	reason object	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	String	Status of the operation. One of: "Success" or "Failure".

表 6-117 StatusDetails v1 meta 数据结构说明

参数	参数类型	描述
causes	Array of StatusCause object	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action - for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described)

表 6-118 StatusCause v1 meta 数据结构说明

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.
reason	StatusCause Type object	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 6-119 StatusCause reason 字段值范围说明

值	描述
FieldValueNotFound	CauseTypeFieldValueNotFound is used to report failure to find a requested value.(e.g. looking up an ID)
FieldValueRequired	CauseTypeFieldValueRequired is used to report required values that are not provided (e.g. empty strings, null values, or empty arrays).
FieldValueDuplicate	CauseTypeFieldValueDuplicate is used to report collisions of values that must be unique (e.g. unique IDs).
FieldValueInvalid	CauseTypeFieldValueInvalid is used to report malformed values (e.g. failed regex match).
FieldValueNotSupported	CauseTypeFieldValueNotSupported is used to report valid (as per formatting rules) values that cannot be handled (e.g. an enumerated string).
UnexpectedServerResponse	CauseTypeUnexpectedServerResponse is used to report when the server responded to the client without the expected return type. The presence of this cause indicates the error may be due to an intervening proxy or the server software malfunctioning.

表 6-120 ListMeta v1 meta 数据结构说明

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only

表 6-121 Status v1 中 reason 值说明

名称	值	描述
StatusReasonUnknown	“ ”	StatusReasonUnknown means the server has declined to indicate a specific reason. The details field may contain other information about this error. Status code 500
StatusReasonUnauthorized	Unauthorized	StatusReasonUnauthorized means the server can be reached and understood the request, but requires the user to present appropriate authorization credentials (identified by the WWW-Authenticate header) in order for the action to be completed. If the user has specified credentials on the request, the server considers them insufficient. Status code 401

名称	值	描述
StatusReasonForbidden	Forbidden	StatusReasonForbidden means the server can be reached and understood the request, but refuses to take any further action. It is the result of the server being configured to deny access for some reason to the requested resource by the client. Details (optional): "kind" string - the kind attribute of the forbidden resource on some operations may differ from the requested resource. "id" string - the identifier of the forbidden resource Status code 403
StatusReasonNotFound	NotFound	StatusReasonNotFound means one or more resources required for this operation could not be found. Details (optional): "kind" string - the kind attribute of the missing resource on some operations may differ from the requested resource. "id" string - the identifier of the missing resource Status code 404
StatusReasonAlreadyExists	AlreadyExists	StatusReasonAlreadyExists means the resource you are creating already exists. Details (optional): "kind" string - the kind attribute of the conflicting resource "id" string - the identifier of the conflicting resource Status code 409
StatusReasonConflict	Conflict	StatusReasonConflict means the requested operation cannot be completed due to a conflict in the operation. The client may need to alter the request. Each resource may define custom details that indicate the nature of the conflict. Status code 409
StatusReasonGone	Gone	StatusReasonGone means the item is no longer available at the server and no forwarding address is known. Status code 410

名称	值	描述
StatusReasonInvalid	Invalid	<p>StatusReasonInvalid means the requested create or update operation cannot be completed due to invalid data provided as part of the request. The client may need to alter the request. When set, the client may use the StatusDetailsmessage field as a summary of the issues encountered.</p> <p>Details (optional):</p> <p>"kind" string - the kind attribute of the invalid resource</p> <p>"id" string - the identifier of the invalid resource</p> <p>"causes" - one or more StatusCause entries indicating the data in the provided resource that was invalid. The code, message, and field attributes will be set.</p> <p>Status code 422</p>
StatusReasonServerTimeout	ServerTimeout	<p>StatusReasonServerTimeout means the server can be reached and understood the request, but cannot complete the action in a reasonable time. The client should retry the request. This is probably due to temporary server load or a transient communication issue with another server. Status code 500 is used because the HTTP spec provides no suitable server-requested client retry and the 5xx class represents actionable errors.</p> <p>Details (optional):</p> <p>"kind" string - the kind attribute of the resource being acted on.</p> <p>"id" string - the operation that is being attempted.</p> <p>"retryAfterSeconds" int32 - the number of seconds before the operation should be retried</p> <p>Status code 500</p>
StatusReasonTimeout	Timeout	<p>StatusReasonTimeout means that the request could not be completed within the given time. Clients can get this response only when they specified a timeout param in the request, or if the server cannot complete the operation within a reasonable amount of time. The request might succeed with an increased value of timeout param. The client <i>should</i> wait at least the number of seconds specified by the retryAfterSeconds field. Details (optional): "retryAfterSeconds" int32 - the number of seconds before the operation should be retried</p> <p>Status code 504</p>

名称	值	描述
StatusReasonBadRequest	BadRequest	StatusReasonBadRequest means that the request itself was invalid, because the request does not make any sense, for example deleting a read-only object. This is different from StatusReasonInvalid above which indicates that the API call could possibly succeed, but the data was invalid. API calls that return BadRequest can never succeed.
StatusReasonMethodNotAllowed	MethodNotAllowed	StatusReasonMethodNotAllowed means that the action the client attempted to perform on the resource was not supported by the code - for instance, attempting to delete a resource that can only be created. API calls that return MethodNotAllowed can never succeed.
StatusReasonInternalError	InternalError	StatusReasonInternalError indicates that an internal error occurred, it is unexpected and the outcome of the call is unknown. Details (optional): "causes" - The original error Status code 500
StatusReasonExpired	Expired	StatusReasonExpired indicates that the request is invalid because the content you are requesting has expired and is no longer available. It is typically associated with watches that cannot be serviced. Status code 410 (gone)
StatusReasonServiceUnavailable	ServiceUnavailable	StatusReasonServiceUnavailable means that the request itself was valid, but the requested service is unavailable at this time. Retrying the request after some time might succeed. Status code 503

6.3 数据结构

表 6-122 PodTemplate 数据结构说明

参数	是否必选	参数类型	描述
kind	Yes	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is PodTemplate .
apiVersion	Yes	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	Yes	ObjectMeta object	-
template	Yes	PodTemplate Spec object	-

表 6-123 Pod 数据结构说明

参数	是否必选	参数类型	描述
kind	Yes	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Pod .

参数	是否必选	参数类型	描述
apiVersion	Yes	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	Yes	ObjectMeta object	-
spec	Yes	podSpec object	-
status	No	PodStatus object	Most recently observed status of the pod.

表 6-124 PodStatus 字段数据结构说明

参数	是否必选	参数类型	描述
phase	No	String	Current condition of the pod. 说明 Pod的状态有： - Pending: means the pod has been accepted by the system, but one or more of the containers has not been started. This includes time before being bound to a node, as well as time spent pulling images onto the host. - Running: the pod has been bound to a node and all of the containers have been started. At least one container is still running or is in the process of being restarted. - Succeeded: all containers in the pod have voluntarily terminated with a container exit code of 0, and the system is not going to restart any of these containers. - Failed: all containers in the pod have terminated, and at least one container has terminated in a failure (exited with a non-zero exit code or was stopped by the system). - Unknown: for some reason the state of the pod could not be obtained, typically due to an error in communicating with the host of the pod.
conditions	No	Array of PodConditions objects	Current service state of the pod.
message	No	String	A human readable message indicating details about why the pod is in this condition.
reason	No	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'OutOfDisk'
hostIP	No	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.

参数	是否必选	参数类型	描述
podIP	No	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
startTime	No	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.
containerStatuses	No	Array of containerStatuses objects	The list has one entry per container in the manifest. Each entry is currently the output of container inspect.
initContainerStatuses	No	Array of containerStatuses objects	The list has one entry per init container in the manifest. The most recent successful init container will have ready = true, the most recently started container will have startTime set.
qosClass	No	String	The Quality of Service (QOS) classification assigned to the pod based on resource requirements. Can be <ul style="list-style-type: none">- Guaranteed- Burstable- BestEffort
podNetworks	No	Array of PodNetworkInterface objects	Complete list of Networks attached to this pod

表 6-125 PodConditions 字段数据结构说明

参数	是否必选	参数类型	描述
type	No	String	Type of the condition. Currently only Ready. Resizing - An user trigger resize of pvc has been started 说明 PodCondition的状态有： - PodScheduled: represents status of the scheduling process for this pod. - Ready: pod is able to service requests and should be added to the load balancing pools of all matching services. - Initialized: all init containers in the pod have started successfully. - Unscheduleable: the scheduler cannot schedule the pod right now, for example due to insufficient resources in the cluster.
status	No	String	Status of the condition. Can be True, False, or Unknown.
lastProbeTime	No	String	Last time we probed the condition.
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.
reason	No	String	Unique, one-word, CamelCase reason for the condition's last transition.
message	No	String	Human-readable message indicating details about last transition.

表 6-126 containerStatuses 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	This must be a DNS_LABEL. Each container in a pod must have a unique name. Cannot be updated.

参数	是否必选	参数类型	描述
state	No	ContainerState object	Details about the container's current condition.
lastState	No	ContainerState object	Details about the container's last termination condition.
ready	No	Boolean	Specifies whether the container has passed its readiness probe.
restartCount	No	Integer	The number of times the container has been restarted, currently based on the number of dead containers that have not yet been removed. Note that this is calculated from dead containers. However, those containers are subject to garbage collection. This value will get capped at 5 by GC.
image	Yes	String	The image the container is running.
imageID	No	String	ID of the container's image.
containerID	No	String	Container's ID in the format 'docker://'.

表 6-127 ContainerState 字段数据结构说明

参数	是否必选	参数类型	描述
waiting	No	ContainerStateWaiting object	Details about a waiting container
running	No	ContainerStateRunning object	Details about a running container
terminated	No	terminated object	Details about a terminated container

表 6-128 ContainerStateWaiting 字段数据结构说明

参数	是否必选	参数类型	描述
reason	No	String	(Brief) Reason the container is not yet running.
message	No	String	Message regarding why the container is not yet running.

表 6-129 ContainerStateRunning 字段数据结构说明

参数	是否必选	参数类型	描述
startedAt	No	String	Time at which the container was last (re-)started.

表 6-130 terminated 字段数据结构说明

参数	是否必选	参数类型	描述
exitCode	No	Integer	Exit status from the last termination of the container.
signal	No	Integer	Signal from the last termination of the container.
reason	No	String	(Brief) reason from the last termination of the container.
message	No	String	Message regarding the last termination of the container.
startedAt	No	String	Time at which previous execution of the container started.
finishedAt	No	String	Time at which the container last terminated.
containerID	No	String	Container's ID in the format 'docker://'

表 6-131 ObjectMeta 字段数据结构说明

参数	是否 必选	参数类型	描述
name	Yes	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. 0 characters < name length ≤ 63 characters. The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
cluster Name	No	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
initializers	No	initializer s object	An initializer is a controller which enforces some system invariant at object creation time. This field is a list of initializers that have not yet acted on this object. If nil or empty, this object has been completely initialized. Otherwise, the object is considered uninitialized and is hidden (in list/watch and get calls) from clients that haven't explicitly asked to observe uninitialized objects. When an object is created, the system will populate this list with the current set of initializers. Only privileged users may set or modify this list. Once it is empty, it may not be modified further by any user.
enable	No	Boolean	Enable identify whether the resource is available.

参数	是否必选	参数类型	描述
generateName	No	String	<p>An optional prefix used by the server to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different from the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409. Instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified.</p> <p>0 characters < generated name length ≤ 253 characters.</p> <p>The generated name must be a regular expression <code>[a-z0-9]([-a-z0-9]*[a-z0-9])?</code>.</p>
namespace	No	String	<p>Namespace defines the space within each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated.</p> <p>0 characters < namespace length ≤ 63 characters.</p> <p>The namespace must be a regular expression <code>[a-z0-9]([-a-z0-9]*[a-z0-9])?</code>.</p>
selfLink	No	String	<p>A URL representing this object. Populated by the system. Read-only.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
uid	No	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>

参数	是否必选	参数类型	描述
resourceVersion	No	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
generation	No	Integer	<p>A sequence number representing a specific generation of the desired state. Currently only implemented by replication controllers. Populated by the system. Read-only.</p>
creationTimestamp	No	String	<p>A timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
deletionTimestamp	No	String	<p>RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource will be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field. Once set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. Once the resource is deleted in the API, the Kubelet will send a hard termination signal to the container. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only.</p>

参数	是否必选	参数类型	描述
deletionGracePeriodSeconds	No	Integer	Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.
labels	No	Object	Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services.
annotations	No	annotations object	An unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects. 说明 不同的资源类型必填的annotation不同，详见具体的资源API。
ownerReferences	No	Array of ownerReferences objects	List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.
finalizers	No	Array of strings	Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed.

表 6-132 annotations 字段数据结构说明

参数	是否必选	参数类型	描述
pod.logcollection.kubernetes.io	No	Array of strings	<p>用于配置需要采集标准输出日志的容器列表。若不配置表示收集所有容器列表。</p> <ul style="list-style-type: none"> 示例1：收集所有 containers。 pod annotation: log.stdoutcollection.kubernetes.io: {"collectionContainers":[]} 示例2：采集容器 container0，其中container0为容器名。 pod annotation: log.stdoutcollection.kubernetes.io: {"collectionContainers": ["container0"]}
paas.storage.io/cryptKeyId	No	String	<p>加密密钥ID。 存储类型为SFS或EVS且需创建加密卷时，才需配置该参数。 密钥ID可从“安全控制台 > 数据加密服务 > 密钥管理”页面获取。</p>
paas.storage.io/cryptAlias	No	String	<p>加密密钥别名。 存储类型为SFS且需创建加密卷时，才需配置该参数。 密钥名称可从“安全控制台 > 数据加密服务 > 密钥管理”页面获取。</p>
paas.storage.io/cryptDomainId	No	String	<p>租户的DomainId 存储类型为SFS且需创建加密卷时，才需配置该参数。</p>

表 6-133 initializers 字段数据结构说明

参数	是否必选	参数类型	描述
pending	No	Array of pending objects	<p>Pending is a list of initializers that must execute in order before this object is visible. When the last pending initializer is removed, and no failing result is set, the initializers struct will be set to nil and the object is considered as initialized and visible to all clients.</p>

参数	是否必选	参数类型	描述
result	No	status object	If result is set with the Failure field, the object will be persisted to storage and then deleted, ensuring that other clients can observe the deletion.

表 6-134 pending 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	name of the process that is responsible for initializing this object.

表 6-135 ownerReferences 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	API version of the referent.
blockOwnerDeletion	No	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Default to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
kind	Yes	String	Kind of the referent.
name	Yes	String	Name of the referent.
uid	No	String	UID of the referent.
controller	No	Boolean	If true, this reference points to the managing controller.

表 6-136 spec 字段数据结构说明

参数	是否必选	参数类型	描述
replicas	No	Integer	The number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Value range: ≥ 0 . Default: 1
minReadySeconds	No	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its containers crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)
template	Yes	PodTemplate Spec object	-
selector	Yes	Object	A label query over pods that should match the Replicas count. If Selector is empty, it is defaulted to the labels present on the Pod template. Label keys and values that must match in order to be controlled by this replication controller, if empty defaulted to labels on Pod template.

表 6-137 status 字段数据结构说明

参数	是否必选	参数类型	描述
replicas	No	Integer	The most recently observed number of replicas.
availableReplicas	No	Integer	The number of available replicas (ready for at least minReadySeconds) for this replication controller.
readyReplicas	No	Integer	The number of ready replicas for this replication controller.
conditions	No	condition object	Represents the latest available observations of a replication controller's current state.

参数	是否必选	参数类型	描述
observedGeneration	No	Integer	Reflects the generation of the most recently observed replication controller.
FullylabeledReplicas	No	Object	-

表 6-138 PodTemplateSpec 字段数据结构说明

参数	是否必选	参数类型	描述
metadata	No	ObjectMeta object	-
spec	Yes	podSpec object	-

表 6-139 condition 字段数据结构说明

参数	是否必选	参数类型	描述
lastTransitionTime	No	String	The last time the condition transitioned from one status to another.
message	No	String	A human readable message indicating details about the transition.
reason	No	String	The reason for the condition's last transition.
status	No	String	Status of the condition, one of True, False, Unknown.
type	No	String	Type of replication controller condition.

表 6-140 podSpec 字段数据结构说明

参数	是否必选	参数类型	描述
volumes	No	Array of volumes objects	List of volumes that can be mounted by containers belonging to the pod.

参数	是否必选	参数类型	描述
affinity	No	affinity object	If specified, the pod's scheduling constraints. 说明 不允许用户设置affinity，默认使用软反亲和
containers	Yes	Array of containers objects	List of containers belonging to the pod. Containers cannot currently be added or removed. There must be at least one container in a pod. Cannot be updated.
restartPolicy	No	String	Restart policy for all containers within the pod. Value: <ul style="list-style-type: none">• Always• OnFailure• Never Default: Always.
priority	No	Integer	Pod优先级，数值越大优先级越高，默认值为0。 取值范围：[-10, 10]
terminationGracePeriodSeconds	No	Integer	Optional duration in seconds the pod needs to terminate gracefully. May be decreased in delete request. Value must be a non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period will be used instead. The grace period is the duration in seconds after the processes running in the pod are sent a termination signal and the time when the processes are forcibly halted with a kill signal. Set this value longer than the expected cleanup time for your process. Defaults to 30 seconds.

参数	是否必选	参数类型	描述
activeDeadlineSeconds	No	Integer	Optional duration in seconds the pod may be active on the node relative to StartTime before the system will actively try to mark it failed and kill associated containers. Value must be a positive integer. Value range of this parameter: > 0.
dnsPolicy	No	String	Set DNS policy for containers within the pod. Value: <ul style="list-style-type: none"> ClusterFirst Default 说明 dnsPolicy不支持设置为Default Default: ClusterFirst.
hostAliases	No	Array of hostAliases objects	HostAliases is an optional list of hosts and IPs that will be injected into the pod's hosts file if specified. This is only valid for non-hostNetwork pods.
serviceAccountName	No	String	Name of the ServiceAccount used to run this pod. 0 characters < service account name length ≤ 253 characters. The service account name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?. 说明 不支持serviceaccount,该字段不支持设置
serviceAccount	No	String	DeprecatedServiceAccount is a deprecated alias for ServiceAccountName. Deprecated: Use serviceAccountName instead. 说明 不支持serviceaccount,该字段不支持设置

参数	是否必选	参数类型	描述
schedulerName	No	String	If specified, the pod will be dispatched by specified scheduler. If not specified, the pod will be dispatched by default scheduler. 说明 不支持设置指定scheduler name
nodeName	No	String	A request to schedule this pod onto a specific node. If it is non-empty, the scheduler simply schedules this pod onto that node, assuming that it fits resource requirements. 0 characters < node name length ≤ 253 characters. The node name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?. 说明 不支持设置nodeName
nodeSelector	No	Object	NodeSelector is a selector which must be true for the pod to fit on a node. Selector which must match a node's labels for the pod to be scheduled on that node. 说明 不支持设置nodeSelector
automountServiceAccountToken	No	Boolean	AutomountServiceAccountToken indicates whether a service account token should be automatically mounted.
hostNetwork	No	Boolean	Host networking requested for this pod. Use the host's network namespace. If this option is set, the ports that will be used must be specified. Default to false. (不支持设置) 说明 不支持使用主机网络

参数	是否必选	参数类型	描述
hostPID	No	Boolean	A flag indicating whether to use the host's pid namespace. Optional: Default to false. 说明 不支持使用主机PID namespaces
hostIPC	No	Boolean	A flag indicating whether to use the host's ipc namespace. Optional: Default to false. 说明 不支持使用主机IPC namespaces
securityContext	No	PodSecurityContext object	SecurityContext holds pod-level security attributes and common container settings. Defaults to empty.
imagePullSecrets	No	Array of imagePullSecrets objects	An optional list of references to secrets in the same namespace to use for pulling any of the images used by this PodSpec. If specified, these secrets will be passed to individual puller implementations for them to use. 说明 使用的镜像是容器镜像服务界面上的“我的镜像”页签中的镜像时必须设置。

参数	是否必选	参数类型	描述
initContainers	No	Array of containers objects	List of initialization containers belonging to the pod. Init containers are executed in order prior to containers being started. If any init container fails, the pod is considered to have failed and is handled according to its restartPolicy. The name for an init container or normal container must be unique among all containers. Init containers may not have Lifecycle actions, Readiness probes, or Liveness probes. The resourceRequirements of an init container are taken into account during scheduling by finding the highest request/limit for each resource type, and then using the max of that value or the sum of the normal containers. Limits are applied to init containers in a similar fashion. Init containers cannot currently be added or removed.
hostname	No	String	Specifies the hostname of the Pod. If not specified, the pod's hostname will be set to a system-defined value.
subdomain	No	String	If specified, the fully qualified Pod hostname will be "<hostname>.<subdomain>.<pod namespace>.svc<cluster domain>". If not specified, the pod will not have a domainname at all.
tolerations	No	tolerations object	If specified, the pod's tolerations. 说明 不支持设置tolerations

参数	是否必选	参数类型	描述
priorityClassName	No	String	If specified, indicates the pod's priority. "SYSTEM" is a special keyword which indicates the highest priority. Any other name must be defined by creating a PriorityClass object with that name. If not specified, the pod priority will be default or zero if there is no default.

表 6-141 volumes 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Volume name. Must be a DNS_LABEL and unique within the pod. 0 characters < volume name length ≤ 63 characters. The volume name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
secret	No	SecretVolumeSource object	Secret represents a secret that should populate this volume.
persistentVolumeClaim	No	PersistentVolumeClaimVolumeSource object	PersistentVolumeClaimVolumeSource represents a reference to a PersistentVolumeClaim in the same namespace.
localDir	No	LocalDirVolumeSource object	LocalDir represents a LocalDir volume that is created by LVM and mounted into the pod
emptyDir	No	emptyDir object	通过LVM方式创建本地卷的Pod实例。

表 6-142 containers 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the container specified as a DNS_LABEL. Each container in a pod must have a unique name (DNS_LABEL). Cannot be updated. 0 characters < container name length ≤ 63 characters. The container name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?. Cannot be updated.
image	Yes	String	容器镜像地址，具体请参见 获取容器镜像地址 。
command	No	Array of strings	Entrypoint array. Not executed within a shell. The container image's entrypoint is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated.
args	No	Array of strings	Arguments to the entrypoint. The container image's cmd is used if this is not provided. Variable references \$(VAR_NAME) are expanded using the container's environment. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$\$(VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Cannot be updated.
workingDir	No	String	Container's working directory. Defaults to Container's default. Defaults to image's default. Cannot be updated.
ports	No	Array of ContainerPort objects	List of ports to expose from the container. Cannot be updated.

参数	是否必选	参数类型	描述
env	No	Array of EnvVar objects	List of environment variables to set in the container. Cannot be updated.
envFrom	No	Array of EnvFromSource objects	List of sources to populate environment variables in the container. The keys defined within a source must be a C_IDENTIFIER. All invalid keys will be reported as an event when the container is starting. When a key exists in multiple sources, the value associated with the last source will take precedence. Values defined by an Env with a duplicate key will take precedence. Cannot be updated.
resources	No	ResourceRequirements object	Minimum resources the volume should have. Cannot be updated.
volumeMounts	No	Array of volumeMounts objects	Pod volumes to mount into the container's filesystem. Cannot be updated.
volumeDevices	No	Array of volumeDevice objects	volumeDevices is the list of block devices to be used by the container. This is an alpha feature and may change in the future.
livenessProbe	No	Probe object	Periodic probe of container liveness. Container will be restarted if the probe fails. Cannot be updated.
readinessProbe	No	Probe object	Periodic probe of container service readiness. Container will be removed from service endpoints if the probe fails. Cannot be updated.
lifecycle	No	lifecycle object	Actions that the management system should take in response to container lifecycle events. Cannot be updated.

参数	是否必选	参数类型	描述
terminationMessagePath	No	String	Path at which the file to which the container's termination message will be written is mounted into the container's filesystem. Message written is intended to be brief final status, such as an assertion failure message. Defaults to /dev/termination-log. Cannot be updated. Cannot be updated.
terminationMessagePolicy	No	String	Indicate how the termination message should be populated. File will use the contents of terminationMessagePath to populate the container status message on both success and failure. FallbackToLogsOnError will use the last chunk of container log output if the termination message file is empty and the container exited with an error. The log output is limited to 2048 bytes or 80 lines, whichever is smaller. Defaults to File. Cannot be updated. 说明 支持的值： - File: default behavior and will set the container status message to the contents of the container's terminationMessagePath when the container exits. - FallbackToLogsOnError: will read the most recent contents of the container logs for the container status message when the container exits with an error and the terminationMessagePath has no contents.
imagePullPolicy	No	String	Image pull policy. Defaults to Always if the :latest tag is specified, or IfNotPresent otherwise. Cannot be updated. Value: <ul style="list-style-type: none"> • Always • Never • IfNotPresent Cannot be updated. 说明 只支持 Always
securityContext	No	security Context object	Security options the pod should run with

参数	是否必选	参数类型	描述
stdin	No	Boolean	A flag indicating whether this container should allocate a buffer for stdin in the container runtime. If this is not set, reads from stdin in the container will always result in EOF. Default is false.
stdinOnce	No	Boolean	A flag indicating whether the container runtime should close the stdin channel after it has been opened by a single attach. When stdin is true, the stdin stream will remain open across multiple attach sessions. If stdinOnce is set to true, stdin is opened on container start, is empty until the first client attaches to stdin, and then remains open and accepts data until the client disconnects, at which time stdin is closed and remains closed until the container is restarted. If this flag is false, a container process that reads from stdin will never receive an EOF. Default is false.
tty	No	Boolean	A flag indicating whether this container should allocate a TTY for itself, also requires 'stdin' to be true. Default is false.

表 6-143 PodSecurityContext 字段数据结构说明

参数	是否必选	参数类型	描述
seLinuxOptions	No	seLinuxOptions object	-
runAsUser	No	Integer	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence for that container. Value length: > 0 characters.

参数	是否必选	参数类型	描述
runAsNonRoot	No	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in SecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
supplementalGroups	No	Array of integers	A list of groups applied to the first process run in each container, in addition to the container's primary GID. If unspecified, no groups will be added to any container.
fsGroup	No	Integer	A special supplemental group that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: The owning GID will be the FSGroup 2. The setgid bit is set (new files created in the volume will be owned by FSGroup) 3. The permission bits are OR'd with rw-rw.

参数	是否必选	参数类型	描述
fsOwner	No	Integer	A special supplemental owner that applies to all containers in a pod. Some volume types allow the Kubelet to change the ownership of that volume to be owned by the pod: <ol style="list-style-type: none">1. The owning UID will be the FSOwner2. The setgid bit is set (new files created in the volume will be owned by FSOwner)3. The permission bits are OR'd with rw----- If unset, the Kubelet will not modify the ownership and permissions of any volume.

表 6-144 imagePullSecrets 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	Name of the referent. 须知 创建工作负载时，使用的镜像是容器镜像服务界面上的“我的镜像”页签中的镜像，那么该参数的值必须设置为“imagepull-secret”。

表 6-145 SecretVolumeSource 字段数据结构说明

参数	是否必选	参数类型	描述
secretName	No	String	Name of a secret in the pod's namespace.

参数	是否必选	参数类型	描述
items	No	items(KeyToPath) object	If unspecified, each key-value pair in the Data field of the referenced Secret will be projected into the volume as a file whose name is the key and content is the value. If specified, the listed keys will be projected into the specified paths, and unlisted keys will not be present. If a key is specified which is not present in the Secret, the volume setup will error. Paths must be relative and may not contain the '..' path or start with '..'.
defaultMode	否	Integer	Optional: mode bits to use on created files by default. Must be a value between 0 and 0777. Defaults to 0644. Directories within the path are not affected by this setting. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.
optional	否	Boolean	Specify whether the Secret or its keys must be defined

表 6-146 PersistentVolumeClaimVolumeSource 字段数据结构说明

参数	是否必选	参数类型	描述
claimName	No	String	Name of a PersistentVolumeClaim in the same namespace as the pod using this volume.
readOnly	No	Boolean	readOnly here will force the ReadOnly setting in VolumeMounts. Value: <ul style="list-style-type: none">• true• false Default: false.

表 6-147 items(KeyToPath)字段数据结构说明

参数	是否必选	参数类型	描述
key	No	String	The key to project.
path	No	String	The relative path of the file to map the key to. May not be an absolute path. May not contain the path element '!'. May not start with the string '!'.
mode	No	Integer	mode bits to use on this file, must be a value between 0 and 0777. If not specified, the volume defaultMode will be used. This might be in conflict with other options that affect the file mode, like fsGroup, and the result can be other mode bits set.

表 6-148 ContainerPort 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	If specified, this must be an IANA_SVC_NAME and unique within the pod. Each named port in a pod must have a unique name. Name for the port that can be referred to by services. 0 characters < name length ≤ 15 characters. The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
hostPort	No	Integer	Number of the port to expose on the host. If specified, this must be a valid port number, 0 < x < 65536. If HostNetwork is specified, this must match ContainerPort. Most containers do not need this. Value range: [1, 65535]. 说明 不支持配置hostPort

参数	是否必选	参数类型	描述
containerPort	No	Integer	Number of the port to expose on the pod's IP address. This must be a valid port number, $0 < x < 65536$. Value range: [1, 65535].
protocol	No	String	Protocol for port. Value: <ul style="list-style-type: none">• TCP• UDP Default: TCP.
hostIP	No	String	What host IP to bind the external port to. 说明 不支持配置hostIP

表 6-149 EnvVar 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the environment variable. Must be a C_IDENTIFIER.
value	No	String	Variable references \$ (VAR_NAME) are expanded using the previous defined environment variables in the container and any service environment variables. If a variable cannot be resolved, the reference in the input string will be unchanged. The \$(VAR_NAME) syntax can be escaped with a double \$\$, for example, \$\$ (VAR_NAME). Escaped references will never be expanded, regardless of whether the variable exists or not. Defaults to "".
valueFrom	No	EnvVarSource object	Source for the environment variable's value. Cannot be used if value is not empty.

表 6-150 ResourceRequirements 字段数据结构说明

参数	是否必选	参数类型	描述
limits	No	Array of ResourceName objects	Maximum amount of compute resources allowed. 说明 limits和requests值必须相等，否则会报错
requests	No	Array of ResourceName objects	Minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. 云容器实例中Pod规格有限制，具体的限制请参见 约束限制 页面的“Pod规格”部分。

表 6-151 ResourceName 字段可选值说明

参数	是否必选	参数类型	描述
storage	No	String	Volume size, in bytes (e.g. 5Gi = 5GiB = 5 * 1024 * 1024 * 1024)
cpu	No	String	CPU size, in cores. (500m = .5 cores)
memory	No	String	Memory size, in bytes. (500Gi = 500GiB = 500 * 1024 * 1024 * 1024)
localdir	No	String	Local Storage for LocalDir, in bytes. (500Gi = 500GiB = 500 * 1024 * 1024 * 1024)
nvidia.com/gpu-tesla-v100-16GB	No	String	NVIDIA GPU resource, the type may change in different environments, in production environment is nvidia.com/gpu-tesla-v100-16GB now. The value must be an integer and not less than 1.

表 6-152 volumeMounts 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	This must match the Name of a Volume. 0 character < name length ≤ 253 characters. The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.
readOnly	No	Boolean	Mounted read-only if true, read-write otherwise (false or unspecified). Value: <ul style="list-style-type: none">• true• false Default: false.
mountPath	No	String	Path within the container at which the volume should be mounted. Value length: > 0 characters.
subPath	No	String	Path within the volume from which the container's volume should be mounted. Defaults to "" (volume's root).

参数	是否必选	参数类型	描述
mountPropagation	No	String	<p>mountPropagation determines how mounts are propagated from the host to container and the other way around.</p> <p>When not set, MountPropagationHostToContainer is used。 This field is alpha in 1.8 and can be reworked or removed in a future release.</p> <p>说明 可设置的值有：</p> <ul style="list-style-type: none">- HostToContainer: means that the volume in a container will receive new mounts from the host or other containers, but filesystems mounted inside the container will not be propagated to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rslave" in Linux terminology).- Bidirectional: means that the volume in a container will receive new mounts from the host or other containers, and its own mounts will be propagated from the container to the host or other containers. Note that this mode is recursively applied to all mounts in the volume ("rshared" in Linux terminology)

参数	是否必选	参数类型	描述
extendPathMode	No	String	Extend the volume path by appending the pod metadata to the path according to specified pattern, which provides a way of directory isolation and helps prevent the writing conflict between different pods. 说明 可设置的值有： - PodUID: Include PodUID in path - PodName: Include Pod full name in path - PodUID/ContainerName: Include Pod UID and container name in path - PodName/ContainerName: Include Pod full name and container name in path

表 6-153 volumeDevice 数据结构说明

参数	参数类型	描述
name	String	name must match the name of a persistentVolumeClaim in the pod
devicePath	String	devicePath is the path inside of the container that the device will be mapped to.

表 6-154 Probe 字段数据结构说明

参数	是否必选	参数类型	描述
exec	No	exec object	Only one option should be specified. Exec specifies the action to take.
initialDelaySeconds	No	Integer	Number of seconds after the container has started before liveness probes are initiated. Value range: ≥ 0 .

参数	是否必选	参数类型	描述
timeoutSeconds	No	Integer	Number of seconds after which the probe times out. Value range: ≥ 0 . Default: 1.
periodSeconds	No	Integer	How often (in seconds) to perform the probe. Default to 10 seconds. Minimum value is 1. Value range: ≥ 0 . Default: 10.
successThreshold	No	Integer	Minimum consecutive successes for the probe to be considered successful after having failed. Defaults to 1. Must be 1 for liveness. Minimum value is 1. Value range: ≥ 0 . Default: 1.
failureThreshold	No	Integer	Minimum consecutive failures for the probe to be considered failed after having succeeded. Defaults to 3. Minimum value is 1. Value range: ≥ 0 . Default: 3.

表 6-155 lifecycle 字段数据结构说明

参数	是否必选	参数类型	描述
postStart	No	Handler object	PostStart is called immediately after a container is created. If the handler fails, the container is terminated and restarted according to its restart policy. Other management of the container blocks until the hook completes.

参数	是否必选	参数类型	描述
preStop	No	Handler object	<p>PreStop is called immediately before a container is terminated. The container is terminated after the handler completes. The reason for termination is passed to the handler.</p> <p>Regardless of the outcome of the handler, the container is eventually terminated.</p> <p>Other management of the container blocks until the hook completes.</p>

表 6-156 securityContext 字段数据结构说明

参数	是否必选	参数类型	描述
capabilities	No	capabilities object	<p>The capabilities to add/drop when running containers.</p> <p>Defaults to the default set of capabilities granted by the container runtime.</p>
privileged	No	Boolean	<p>Run container in privileged mode. Processes in privileged containers are essentially equivalent to root on the host.</p> <p>Value:</p> <ul style="list-style-type: none"> • true • false <p>Default: false.</p> <p>说明 不允许设置为True</p>

参数	是否必选	参数类型	描述
seLinuxOptions	No	seLinuxOptions object	The SELinux context to be applied to the container. If unspecified, the container runtime will allocate a random SELinux context for each container. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsUser	No	Integer	The UID to run the entrypoint of the container process. Defaults to user specified in image metadata if unspecified. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence.
runAsNonRoot	No	Boolean	Indicates that the container must run as a non-root user. If true, the Kubelet will validate the image at runtime to ensure that it does not run as UID 0 (root) and fail to start the container if it does. If unset or false, no such validation will be performed. May also be set in PodSecurityContext. If set in both SecurityContext and PodSecurityContext, the value specified in SecurityContext takes precedence. Value: <ul style="list-style-type: none">• true• false
readOnlyRootFilesystem	No	Boolean	Whether this container has a read-only root filesystem. Default is false.

参数	是否必选	参数类型	描述
allowPrivilegeEscalation	No	Boolean	AllowPrivilegeEscalation controls whether a process can gain more privileges than its parent process. This bool directly controls if the no_new_privs flag will be set on the container process. AllowPrivilegeEscalation is true always when the container is: 1) run as Privileged 2) has CAP_SYS_ADMIN

表 6-157 seLinuxOptions 字段数据结构说明

参数	是否必选	参数类型	描述
user	No	String	SELinux user label that applies to the container.
role	No	String	SELinux role label that applies to the container.
type	No	String	SELinux type label that applies to the container.
level	No	String	SELinux level label that applies to the container.

表 6-158 items 字段数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Relative path name of the file to be created. Must not be absolute or contain the '..' path. Must be utf-8 encoded. The first item of the relative path must not start with '..'
fieldRef	No	ObjectFieldSelector object	-

参数	是否必选	参数类型	描述
resourceFieldRef	No	ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests. (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 6-159 EnvVarSource 字段数据结构说明

参数	是否必选	参数类型	描述
fieldRef	No	ObjectFieldSelector object	Selects a field of the pod: supports metadata.name, metadata.namespace, metadata.labels, metadata.annotations, spec.nodeName, spec.serviceAccountName, status.hostIP, status.podIP.
resourceFieldRef	No	ResourceFieldSelector object	Selects a resource of the container: only resources limits and requests. (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.
configMapKeyRef	No	ConfigMapKeySelector object	Selects a key of a ConfigMap.
secretKeyRef	No	SecretKeySelector object	Selects a key of a secret in the pod's namespace
processResourceFieldRef	No	ProcessResourceFieldSelector object	Selects a resource of the process: only resources limits and requests (limits.cpu, limits.memory, requests.cpu and requests.memory) are currently supported.

表 6-160 exec 字段数据结构说明

参数	是否必选	参数类型	描述
command	No	Array of strings	Command is the command line to execute inside the container, the working directory for the command is root ('/') in the container's filesystem. The command is simply executed, it is not run inside a shell, so traditional shell instructions (' ', etc) do not work. To use a shell, you need to explicitly call out to that shell. Exit status of 0 is treated as live/healthy and non-zero is unhealthy.

表 6-161 Handler 数据结构说明

参数	是否必选	参数类型	描述
exec	No	exec object	Only one option should be specified. Exec specifies the action to take.

表 6-162 capabilities 字段数据结构说明

参数	是否必选	参数类型	描述
add	No	Array of strings	Added capabilities.
drop	No	Array of strings	Removed capabilities.

表 6-163 ObjectFieldSelector 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	No	String	Version of the schema the FieldPath is written in terms of. Defaults to "v1".
fieldPath	No	String	Path of the field to select in the specified API version.

表 6-164 ResourceFieldSelector 字段数据结构说明

参数	是否必选	参数类型	描述
containerName	No	String	Container name: required for volumes, optional for env vars.
resource	Yes	String	Required: resource to select.
divisor	No	String	Specifies the output format of the exposed resources, defaults to "1".

表 6-165 ConfigMapKeySelector 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	The ConfigMap name to select from
key	No	String	The key to select
optional	No	String	Specify whether the ConfigMap or its key must be defined

表 6-166 SecretKeySelector 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	The secret name to select from
key	No	String	The key to select
optional	No	String	Specify whether the Secret or its key must be defined

表 6-167 ProcessResourceFieldSelector 字段数据结构说明

参数	是否必选	参数类型	描述
processName	No	String	Process name: required for volumes, optional for env vars
resource	Yes	String	Required: resource to select.
divisor	No	Integer	Specifies the output format of the exposed resources, defaults to "1".

表 6-168 EnvFromSource 字段数据结构说明

参数	是否必选	参数类型	描述
prefix	No	String	An optional identifier to prepend to each key in the ConfigMap. Must be a C_IDENTIFIER.
configMapRef	No	ConfigMapEnvSource object	The ConfigMap to select from
secretRef	No	SecretEnvSource object	The Secret to select from

表 6-169 ConfigMapEnvSource 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	The ConfigMap to select from
optional	No	String	Specify whether the ConfigMap must be defined

表 6-170 SecretEnvSource 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	The Secret name to select from
optional	No	String	Specify whether the Secret must be defined

表 6-171 add 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of the resource.
namespaced	No	Boolean	A flag indicating whether a resource is namespaced or not. Default: false.
kind	No	String	kind of the resource.

表 6-172 affinity 字段数据结构说明

参数	是否必选	参数类型	描述
nodeAffinity	No	nodeAffinity object	Describes node affinity scheduling rules for the pod.
podAffinity	No	podAffinity object	Describes pod affinity scheduling rules (e.g. co-locate this pod in the same node, zone, etc. as some other pod(s)).
podAntiAffinity	No	podAffinity object	Describes pod anti-affinity scheduling rules (e.g. avoid putting this pod in the same node, zone, etc. as some other pod(s)).

表 6-173 nodeAffinity 字段数据结构说明

参数	是否必选	参数类型	描述
preferredDuringSchedulingIgnoredDuringExecution	No	preferredDuringSchedulingIgnoredDuringExecution object	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, requiredDuringScheduling affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node matches the corresponding matchExpressions; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
<code>requiredDuringSchedulingIgnoredDuringExecution</code>	No	<code>requiredDuringSchedulingIgnoredDuringExecution</code> object	If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to an update), the system may or may not try to eventually evict the pod from its node.

表 6-174 podAffinity 字段数据结构说明

参数	是否必选	参数类型	描述
<code>preferredDuringSchedulingIgnoredDuringExecution</code>	No	<code>preferredDuringSchedulingIgnoredDuringExecution</code> object	The scheduler will prefer to schedule pods to nodes that satisfy the affinity expressions specified by this field, but it may choose a node that violates one or more of the expressions. The node that is most preferred is the one with the greatest sum of weights, i.e. for each node that meets all of the scheduling requirements (resource request, <code>requiredDuringScheduling</code> affinity expressions, etc.), compute a sum by iterating through the elements of this field and adding "weight" to the sum if the node has pods which match the corresponding <code>podAffinityTerm</code> ; the node(s) with the highest sum are the most preferred.

参数	是否必选	参数类型	描述
requiredDuringSchedulingIgnoredDuringExecution	No	podAffinityTerm object	<p>NOT YET IMPLEMENTED. TODO: Uncomment field once it is implemented. If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system will try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.</p> <p>RequiredDuringSchedulingRequiredDuringExecution []PodAffinityTerm json:"requiredDuringSchedulingRequiredDuringExecution,omitempty" If the affinity requirements specified by this field are not met at scheduling time, the pod will not be scheduled onto the node. If the affinity requirements specified by this field cease to be met at some point during pod execution (e.g. due to a pod label update), the system may or may not try to eventually evict the pod from its node. When there are multiple elements, the lists of nodes corresponding to each podAffinityTerm are intersected, i.e. all terms must be satisfied.</p>

表 6-175 preferredDuringSchedulingIgnoredDuringExecution 字段数据结构说明

参数	是否必选	参数类型	描述
preference	No	preference object	A node selector term, associated with the corresponding weight.
weight	No	Integer	Weight associated with matching the corresponding nodeSelectorTerm, in the range 1-100.

表 6-176 requiredDuringSchedulingIgnoredDuringExecution 字段数据结构说明

参数	是否必选	参数类型	描述
nodeSelectorTerms	No	preference object	Required. A list of node selector terms. The terms are ORed.

表 6-177 preference 字段数据结构说明

参数	是否必选	参数类型	描述
matchExpressions	No	matchExpressions object	Required. A list of node selector requirements. The requirements are ANDed.

表 6-178 matchExpressions 字段数据结构说明

参数	是否必选	参数类型	描述
key	No	String	The label key that the selector applies to.
operator	No	String	Represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists, DoesNotExist. Gt, and Lt.

参数	是否必选	参数类型	描述
values	No	String	An array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. If the operator is Gt or Lt, the values array must have a single element, which will be interpreted as an integer. This array is replaced during a strategic merge patch.

表 6-179 preferredDuringSchedulingIgnoredDuringExecution 字段数据结构说明

参数	是否必选	参数类型	描述
podAffinityTerm	No	podAffinityTerm object	Required. A pod affinity term, associated with the corresponding weight.
weight	No	Integer	weight associated with matching the corresponding podAffinityTerm, in the range 1-100.

表 6-180 podAffinityTerm 字段数据结构说明

参数	是否必选	参数类型	描述
labelSelector	No	labelSelector object	A label query over a set of resources, in this case pods.
namespaces	No	Array of strings	namespaces specifies which namespaces the labelSelector applies to (matches against); null or empty list means "this pod's namespace"

参数	是否必选	参数类型	描述
topologyKey	No	String	This pod should be co-located (affinity) or not co-located (anti-affinity) with the pods matching the labelSelector in the specified namespaces, where co-located is defined as running on a node whose value of the label with key topologyKey matches that of any node on which any of the selected pods is running. For PreferredDuringScheduling pod anti-affinity, empty topologyKey is interpreted as "all topologies" ("all topologies" here means all the topologyKeys indicated by scheduler command-line argument --failure-domains); for affinity and for RequiredDuringScheduling pod anti-affinity, empty topologyKey is not allowed.

表 6-181 labelSelector 字段数据结构说明

参数	是否必选	参数类型	描述
matchExpressions	No	Array of LabelSelectorRequirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	No	Object	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 6-182 LabelSelectorRequirement 字段数据结构说明

参数	是否必选	参数类型	描述
key	No	String	key is the label key that the selector applies to.
operator	No	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.
values	No	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 6-183 hostAliases 字段数据结构说明

参数	是否必选	参数类型	描述
hostnames	No	Array of strings	Hostnames for the above IP address.
ip	No	String	IP address of the host file entry.

表 6-184 tolerations 字段数据结构说明

参数	是否必选	参数类型	描述
effect	No	String	Effect indicates the taint effect to match. Empty means match all taint effects. When specified, allowed values are NoSchedule, PreferNoSchedule and NoExecute.
key	No	String	Key is the taint key that the toleration applies to. Empty means match all taint keys. If the key is empty, operator must be Exists; this combination means to match all values and all keys.

参数	是否必选	参数类型	描述
operator	No	String	Operator represents a key's relationship to the value. Valid operators are Exists and Equal. Defaults to Equal. Exists is equivalent to wildcard for value, so that a pod can tolerate all taints of a particular category.
tolerationSeconds	No	Integer	TolerationSeconds represents the period of time the toleration (which must be of effect NoExecute, otherwise this field is ignored) tolerates the taint. By default, it is not set, which means tolerate the taint forever (do not evict). Zero and negative values will be treated as 0 (evict immediately) by the system.
value	No	String	Value is the taint value the toleration matches to. If the operator is Exists, the value should be empty, otherwise just a regular string.

表 6-185 DeleteOptions 数据结构说明

参数	是否必选	参数类型	描述
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Namespace .

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
gracePeriodSeconds	No	Integer	The duration in seconds before the object should be deleted. Value must be a non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately. Value range of this parameter: > 0.
preconditions	No	preconditions object	Must be fulfilled before a deletion is carried out. If not possible, a 409 Conflict status will be returned.
orphanDependents	No	Boolean	Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list.

参数	是否必选	参数类型	描述
propagationPolicy	No	String	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. 须知 Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.

表 6-186 preconditions 字段数据结构说明

参数	是否必选	参数类型	描述
uid	No	String	Specifies the target UID.

表 6-187 PodNetworkInterface 数据结构说明

参数	参数类型	描述
name	String	Name of the interface inside the pod
network	String	Name of the attached network
ip	Array of strings	IP address(both v4 and v6) of this interface

表 6-188 v1.PodList 数据结构参数说明

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadataString	ListMeta object	-
items	Array of Pod objects	List of pods.

表 6-189 v1.PodTemplateList 数据结构参数说明

参数	参数类型	描述
kind	String	A string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	Versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	ListMeta object	-
items	Array of PodTemplate objects	List of pod templates.

表 6-190 status 字段数据结构说明

参数	参数类型	描述
phase	String	Current condition of the pod.
conditions	PodConditions object	Current service state of the pod.

参数	参数类型	描述
message	String	A human readable message indicating details about why the pod is in this condition.
reason	String	A brief CamelCase message indicating details about why the pod is in this state. e.g. 'OutOfDisk'
hostIP	String	IP address of the host to which the pod is assigned. Empty if not yet scheduled.
podIP	String	IP address allocated to the pod. Routable at least within the cluster. Empty if not yet allocated.
startTime	String	RFC 3339 date and time at which the object was acknowledged by the Kubelet. This is before the Kubelet pulled the container image(s) for the pod.
containerStatuses	containerStatuses object	The list has one entry per container in the manifest. Each entry is currently the output of container inspect.

表 6-191 metadata 字段数据结构说明

参数	参数类型	描述
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only.
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only.

表 6-192 objectReference 字段数据结构说明

参数	参数类型	描述
kind	String	Kind of the referent.
namespace	String	Namespace of the referent.
name	String	Name of the referent.
uid	String	UID of the referent.
apiVersion	String	API version of the referent.
resourceVersion	String	Specific resourceVersion to which this reference is made, if any.
fieldPath	String	Path of the field to select in the specified API version.

表 6-193 status 字段数据结构说明

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	ListMeta object	Standard list metadata.
status	String	Status of the operation. One of: "Success" or "Failure".
message	String	A human-readable description of the status of this operation.
reason	reason object	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
details	StatusDetails object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.

参数	参数类型	描述
code	Integer	Suggested HTTP return code for this status, 0 if not set.

表 6-194 StatusDetails 数据结构说明

参数	参数类型	描述
causes	Array of StatusCause objects	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	String	The group attribute of the resource associated with the status StatusReason.
kind	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind
name	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	Integer	If specified, the time in seconds before the operation should be retried. Some errors may indicate the client must take an alternate action - for those errors this field may indicate how long to wait before taking the alternate action.
uid	String	UID of the resource. (when there is a single resource which can be described)

表 6-195 StatusCause 数据结构说明

参数	参数类型	描述
field	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.

参数	参数类型	描述
reason	StatusCauseReason object	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 6-196 StatusCause reason 字段值范围说明

参数	描述
FieldValueNotFound	CauseTypeFieldValueNotFound is used to report failure to find a requested value.(e.g. looking up an ID)
FieldValueRequired	CauseTypeFieldValueRequired is used to report required values that are not provided (e.g. empty strings, null values, or empty arrays).
FieldValueDuplicate	CauseTypeFieldValueDuplicate is used to report collisions of values that must be unique (e.g. unique IDs).
FieldValueInvalid	CauseTypeFieldValueInvalid is used to report malformed values (e.g. failed regex match).
FieldValueNotSupported	CauseTypeFieldValueNotSupported is used to report valid (as per formatting rules) values that cannot be handled (e.g. an enumerated string).
UnexpectedServerResponse	CauseTypeUnexpectedServerResponse is used to report when the server responded to the client without the expected return type. The presence of this cause indicates the error may be due to an intervening proxy or the server software malfunctioning.

表 6-197 ListMeta 数据结构说明

参数	参数类型	描述
continue	String	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response

参数	参数类型	描述
resourceVersion	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only
selfLink	String	SelfLink is a URL representing this object. Populated by the system. Read-only

表 6-198 reason

名称	值	描述
StatusReasonUnknown	“ ”	StatusReasonUnknown means the server has declined to indicate a specific reason. The details field may contain other information about this error. Status code 500
StatusReasonUnauthorized	Unauthorized	StatusReasonUnauthorized means the server can be reached and understood the request, but requires the user to present appropriate authorization credentials (identified by the WWW-Authenticate header) in order for the action to be completed. If the user has specified credentials on the request, the server considers them insufficient. Status code 401
StatusReasonForbidden	Forbidden	StatusReasonForbidden means the server can be reached and understood the request, but refuses to take any further action. It is the result of the server being configured to deny access for some reason to the requested resource by the client. Details (optional): "kind" string - the kind attribute of the forbidden resource on some operations may differ from the requested resource. "id" string - the identifier of the forbidden resource Status code 403

名称	值	描述
StatusReasonNotFound	NotFound	StatusReasonNotFound means one or more resources required for this operation could not be found. Details (optional): "kind" string - the kind attribute of the missing resource on some operations may differ from the requested resource. "id" string - the identifier of the missing resource Status code 404
StatusReasonAlreadyExists	AlreadyExists	StatusReasonAlreadyExists means the resource you are creating already exists. Details (optional): "kind" string - the kind attribute of the conflicting resource "id" string - the identifier of the conflicting resource Status code 409
StatusReasonConflict	Conflict	StatusReasonConflict means the requested operation cannot be completed due to a conflict in the operation. The client may need to alter the request. Each resource may define custom details that indicate the nature of the conflict. Status code 409
StatusReasonGone	Gone	StatusReasonGone means the item is no longer available at the server and no forwarding address is known. Status code 410
StatusReasonInvalid	Invalid	StatusReasonInvalid means the requested create or update operation cannot be completed due to invalid data provided as part of the request. The client may need to alter the request. When set, the client may use the StatusDetailsmessage field as a summary of the issues encountered. Details (optional): "kind" string - the kind attribute of the invalid resource "id" string - the identifier of the invalid resource "causes" - one or more StatusCause entries indicating the data in the provided resource that was invalid. The code, message, and field attributes will be set. Status code 422

名称	值	描述
StatusReasonServerTimeout	ServerTimeout	<p>StatusReasonServerTimeout means the server can be reached and understood the request, but cannot complete the action in a reasonable time. The client should retry the request. This is probably due to temporary server load or a transient communication issue with another server. Status code 500 is used because the HTTP spec provides no suitable server-requested client retry and the 5xx class represents actionable errors.</p> <p>Details (optional):</p> <p>"kind" string - the kind attribute of the resource being acted on.</p> <p>"id" string - the operation that is being attempted.</p> <p>"retryAfterSeconds" Integer - the number of seconds before the operation should be retried</p> <p>Status code 500</p>
StatusReasonTimeout	Timeout	<p>StatusReasonTimeout means that the request could not be completed within the given time. Clients can get this response only when they specified a timeout param in the request, or if the server cannot complete the operation within a reasonable amount of time. The request might succeed with an increased value of timeout param. The client <i>should</i> wait at least the number of seconds specified by the retryAfterSeconds field. Details (optional): "retryAfterSeconds" int32 - the number of seconds before the operation should be retried</p> <p>Status code 504</p>
StatusReasonTooManyRequests	TooManyRequests	<p>StatusReasonTooManyRequests means the server experienced too many requests within a given window and that the client must wait to perform the action again. A client may always retry the request that led to this error, although the client should wait at least the number of seconds specified by the retryAfterSeconds field.</p> <p>Details (optional):</p> <p>"retryAfterSeconds" int32 - the number of seconds before the operation should be retried</p> <p>Status code 429</p>

名称	值	描述
StatusReasonBadRequest	BadRequest	StatusReasonBadRequest means that the request itself was invalid, because the request does not make any sense, for example deleting a read-only object. This is different from StatusReasonInvalid above which indicates that the API call could possibly succeed, but the data was invalid. API calls that return BadRequest can never succeed.
StatusReasonMethodNotAllowed	MethodNotAllowed	StatusReasonMethodNotAllowed means that the action the client attempted to perform on the resource was not supported by the code - for instance, attempting to delete a resource that can only be created. API calls that return MethodNotAllowed can never succeed.
StatusReasonInternalError	InternalError	StatusReasonInternalError indicates that an internal error occurred, it is unexpected and the outcome of the call is unknown. Details (optional): "causes" - The original error Status code 500
StatusReasonExpired	Expired	StatusReasonExpired indicates that the request is invalid because the content you are requesting has expired and is no longer available. It is typically associated with watches that cannot be serviced. Status code 410 (gone)
StatusReasonServiceUnavailable	ServiceUnavailable	StatusReasonServiceUnavailable means that the request itself was valid, but the requested service is unavailable at this time. Retrying the request after some time might succeed. Status code 503

表 6-199 WatchEvent 数据结构说明

参数	参数类型	描述
type	String	Type of Event. Can be - Added - Modified - Deleted - Error

参数	参数类型	描述
Object	String	Object is: - If Type is Added or Modified: the new state of the object. - If Type is Deleted: the state of the object immediately before deletion. - If Type is Error: Status is recommended; - other types may make sense depending on context.

表 6-200 Deployment 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	Yes	ObjectMeta object	Standard object metadata.
spec	Yes	DeploymentSpec object	Specification of the desired behavior of the Deployment.
status	No	DeploymentStatus object	Most recently observed status of the Deployment.

表 6-201 DeploymentSpec 字段数据结构说明

参数	是否必选	参数类型	描述
minReadySeconds	No	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its containers crashing, for it to be considered available. Defaults to 0 (pod will be considered available as soon as it is ready)

参数	是否必选	参数类型	描述
paused	No	Boolean	Indicates that the deployment is paused.
progressDeadlineSeconds	No	Integer	The maximum time in seconds for a deployment to make progress before it is considered to be failed. The deployment controller will continue to process failed deployments and a condition with a ProgressDeadlineExceeded reason will be surfaced in the deployment status. Once autoRollback is implemented, the deployment controller will automatically rollback failed deployments. Note that progress will not be estimated during the time a deployment is paused. Defaults to 600s.
replicas	No	Integer	Number of desired pods. This is a pointer to distinguish between explicit zero and not specified. Defaults to 1. replicas为1表示单个Pod，没有高可用能力，建议设置为多个。
priority	No	Integer	负载优先级，数值越大优先级越高，默认值为0。 取值范围：[-10, 10]
revisionHistoryLimit	No	Integer	The number of old ReplicaSets to retain to allow rollback. This is a pointer to distinguish between explicit zero and not specified. Defaults to 2.
selector	No	labelSelector object	Label selector for pods. Existing ReplicaSets whose pods are selected by this will be the ones affected by this deployment.
strategy	No	DeploymentStrategy object	The deployment strategy to use to replace existing pods with new ones.
template	Yes	PodTemplateSpec object	Template describes the pods that will be created.

表 6-202 DeploymentStatus 字段数据结构说明

参数	是否必选	参数类型	描述
availableReplicas	No	Integer	Total number of available pods (ready for at least minReadySeconds) targeted by this deployment.
collisionCount	No	Integer	Count of hash collisions for the Deployment. The Deployment controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ReplicaSet.
conditions	No	Array of DeploymentCondition objects	Represents the latest available observations of a deployment's current state.
observedGeneration	No	Integer	The generation observed by the deployment controller
readyReplicas	No	Integer	Total number of ready pods targeted by this deployment
replicas	No	Integer	Total number of non-terminated pods targeted by this deployment (their labels match the selector).
unavailableReplicas	No	Integer	Total number of unavailable pods targeted by this deployment.
updatedReplicas	No	Integer	Total number of non-terminated pods targeted by this deployment that have the desired template spec.

表 6-203 DeploymentStrategy 字段数据结构说明

参数	是否必选	参数类型	描述
rollingUpdate	Yes	RollingUpdateDeployment object	Rolling update config params. Present only if DeploymentStrategyType = RollingUpdate.
type	No	String	Type of deployment. Can be "Recreate" or "RollingUpdate". Default is RollingUpdate.

表 6-204 DeploymentCondition 字段数据结构说明

参数	是否必选	参数类型	描述
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.
lastUpdateTime	No	String	The last time this condition was updated.
message	No	String	A human readable message indicating details about the transition.
reason	No	String	The reason for the condition's last transition.
status	No	String	Status of the condition, one of True, False, Unknown.
type	No	String	Type of deployment condition. Can be "Available", "Progressing", "ReplicaFailure"

表 6-205 RollingUpdateDeployment 字段数据结构说明

参数	是否必选	参数类型	描述
maxSurge	No	Integer	The maximum number of pods that can be scheduled above the desired number of pods. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). This cannot be 0 if MaxUnavailable is 0. Absolute number is calculated from percentage by rounding up. Defaults to 25%. Example: when this is set to 30%, the new RC can be scaled up immediately when the rolling update starts, such that the total number of old and new pods do not exceed 130% of desired pods. Once old pods have been killed, new RC can be scaled up further, ensuring that total number of pods running at any time during the update is at most 130% of desired pods.

参数	是否必选	参数类型	描述
maxUnavailable	No	Integer	The maximum number of pods that can be unavailable during the update. Value can be an absolute number (ex: 5) or a percentage of desired pods (ex: 10%). Absolute number is calculated from percentage by rounding down. This cannot be 0 if MaxSurge is 0. Defaults to 25%. Example: when this is set to 30%, the old RC can be scaled down to 70% of desired pods immediately when the rolling update starts. Once new pods are ready, old RC can be scaled down further, followed by scaling up the new RC, ensuring that the total number of pods available at all times during the update is at least 70% of desired pods.

表 6-206 DeploymentList v1 apps 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	No	ListMeta object	Standard object metadata.
items	Yes	Array of Deployment objects	Items is the list of Deployments

表 6-207 StatefulSet 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	Yes	ObjectMeta object	Standard list metadata.
spec	Yes	StatefulSetSpec object	Spec defines the desired identities of pods in this set.
status	No	StatefulSetStatus object	Status is the current status of Pods in this StatefulSet. This data may be out of date by some window of time.

表 6-208 StatefulSetStatus 字段数据结构说明

参数	是否必选	参数类型	描述
observedGeneration	No	Integer	Most recent generation observed by this autoscaler.
replicas	No	Integer	Replicas is the number of actual replicas.
currentReplicas	No	Integer	currentReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by currentRevision.
currentRevision	No	String	currentRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [0,currentReplicas).

参数	是否必选	参数类型	描述
readyReplicas	No	Integer	readyReplicas is the number of Pods created by the StatefulSet controller that have a Ready Condition.
updateRevision	No	String	updateRevision, if not empty, indicates the version of the StatefulSet used to generate Pods in the sequence [replicas-updatedReplicas,replicas)
updatedReplicas	No	Integer	updatedReplicas is the number of Pods created by the StatefulSet controller from the StatefulSet version indicated by updateRevision.
collisionCount	No	Integer	collisionCount is the count of hash collisions for the StatefulSet. The StatefulSet controller uses this field as a collision avoidance mechanism when it needs to create the name for the newest ControllerRevision
conditions	No	Array of StatefulSetCondition objects	Represents the latest available observations of a statefulset's current state.

表 6-209 StatefulSetSpec 字段数据结构说明

参数	是否必选	参数类型	描述
replicas	No	Integer	Replicas is the desired number of replicas of the given Template. These are replicas in the sense that they are instantiations of the same Template, but individual replicas also have a consistent identity. If unspecified, defaults to 1. replicas为1表示单个Pod，没有高可用能力，建议设置为多个。
priority	No	Integer	负载优先级，数值越大优先级越高，默认值为0。 取值范围：[-10, 10]

参数	是否必选	参数类型	描述
podManagementPolicy	No	String	<p>podManagementPolicy controls how pods are created during initial scale up, when replacing pods on nodes, or when scaling down.</p> <p>can be: OrderedReady, Parallel</p> <p>The default policy is OrderedReady, where pods are created in increasing order (pod-0, then pod-1, etc) and the controller will wait until each pod is ready before continuing. When scaling down, the pods are removed in the opposite order.</p> <p>The alternative policy is Parallel which will create pods in parallel to match the desired scale without waiting, and on scale down will delete all pods at once.</p>
revisionHistoryLimit	No	Integer	<p>revisionHistoryLimit is the maximum number of revisions that will be maintained in the StatefulSet's revision history. The revision history consists of all revisions not represented by a currently applied StatefulSetSpec version. The default value is 10.</p>
updateStrategy	No	StatefulSetUpdateStrategy object	<p>updateStrategy indicates the StatefulSetUpdateStrategy that will be employed to update Pods in the StatefulSet when a revision is made to Template.</p>
serviceName	Yes	String	<p>ServiceName is the name of the service that governs this StatefulSet. This service must exist before the StatefulSet, and is responsible for the network identity of the set. Pods get DNS/hostnames that follow the pattern: pod-specific-string.serviceName.default.svc.cluster.local where "pod-specific-string" is managed by the StatefulSet controller.</p>

参数	是否必选	参数类型	描述
volumeClaimTemplates	No	PersistentVolumeClaim object	VolumeClaimTemplates is a list of claims that pods are allowed to reference. The StatefulSet controller is responsible for mapping network identities to claims in a way that maintains the identity of a pod. Every claim in this list must have at least one matching (by name) volumeMount in one container in the template. A claim in this list takes precedence over any volumes in the template, with the same name. 注意：当前仅支持挂载EVS云硬盘卷
selector	Yes	labelSelector object	Selector is a label query over pods that should match the replica count. If empty, defaulted to labels on the pod template.
template	Yes	PodTemplateSpec object	Template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.

表 6-210 StatefulSetUpdateStrategy 字段数据结构说明

参数	是否必选	参数类型	描述
rollingUpdate	No	RollingUpdateStatefulSetStrategy object	RollingUpdate is used to communicate parameters when Type is RollingUpdateStatefulSetStrategyType.

参数	是否必选	参数类型	描述
type	No	String	Type indicates the type of the StatefulSetUpdateStrategy. can be: - RollingUpdate : indicates that update will be applied to all Pods in the StatefulSet with respect to the StatefulSet ordering constraints. When a scale operation is performed with this strategy, new Pods will be created from the specification version indicated by the StatefulSet's updateRevision. - OnDelete : triggers the legacy behavior. Version tracking and ordered rolling restarts are disabled. Pods are recreated from the StatefulSetSpec when they are manually deleted. When a scale operation is performed with this strategy, specification version indicated by the StatefulSet's currentRevision.

表 6-211 RollingUpdateStatefulSetStrategy 字段数据结构说明

参数	是否必选	参数类型	描述
partition	No	Integer	Partition indicates the ordinal at which the StatefulSet should be partitioned.

表 6-212 StatefulSetCondition 字段数据结构说明

参数	是否必选	参数类型	描述
type	No	String	Type of the condition. Currently only Ready.
status	No	String	Status of the condition. Can be True, False, or Unknown.
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.

参数	是否必选	参数类型	描述
reason	No	String	Unique, one-word, CamelCase reason for the condition's last transition.
message	No	String	Human-readable message indicating details about last transition.

表 6-213 PersistentVolumeClaim 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	ObjectMeta object	Standard object's metadata.
spec	PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requested by a pod author.
status	PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only.

表 6-214 PersistentVolumeClaimStatus 字段数据结构说明

参数	是否必选	参数类型	描述
accessModes	No	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. ReadWriteOnce - can be mounted read/write mode to exactly 1 host ReadOnlyMany - can be mounted in read-only mode to many hosts ReadWriteMany - can be mounted in read/write mode to many hosts
capacity	No	Array of ResourceName objects	Represents the actual resources of the underlying volume.
phase	No	String	Phase represents the current phase of PersistentVolumeClaim. pending - used for PersistentVolumeClaims that are not yet bound Bound - used for PersistentVolumeClaims that are bound Lost - used for PersistentVolumeClaims that lost their underlying PersistentVolume. The claim was bound to a PersistentVolume and this volume does not exist any longer and all data on it was lost.
conditions	No	Array of PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to ResizeStarted.

表 6-215 PersistentVolumeClaimCondition 字段数据结构说明

参数	是否必选	参数类型	描述
type	No	String	Type of the condition. Currently only Ready. Resizing - An user trigger resize of pvc has been started
status	No	String	Status of the condition. Can be True, False, or Unknown.
lastProbeTime	No	String	Last time we probed the condition.
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.
reason	No	String	Unique, one-word, CamelCase reason for the condition's last transition.
message	No	String	Human-readable message indicating details about last transition.

表 6-216 PersistentVolumeClaimSpec 字段数据结构说明

参数	是否必选	参数类型	描述
volumeName	No	String	VolumeName is the binding reference to the PersistentVolume backing this claim.
accessModes	Yes	Array of strings	AccessModes contains the desired access modes the volume should have. ReadWriteOnce – the volume can be mounted as read-write by a single node ReadOnlyMany – the volume can be mounted read-only by many nodes ReadWriteMany – the volume can be mounted as read-write by many nodes
resources	Yes	ResourceRequirements object	Resources represents the minimum resources the volume should have.

参数	是否必选	参数类型	描述
selector	No	labelSelector object	A label query over volumes to consider for binding.
storageClassName	Yes	String	Name of the StorageClass required by the claim. 支持如下字段： <ul style="list-style-type: none"> 云硬盘存储卷（EVS） 当前支持高I/O（sas）、超高I/O（ssd）和普通I/O（sata） 文件存储卷（SFS） 当前支持标准文件协议类型（nfs-rw）
volumeMode	No	String	volumeMode defines what type of volume is required by the claim. Can be: <ul style="list-style-type: none"> - Block: the volume will not be formatted with a filesystem and will remain a raw block device - Filesystem: the volume will be or is formatted with a filesystem

表 6-217 StatefulSetList v1 apps 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
-	No	ListMeta object	-
items	Yes	Array of StatefulSet objects	Items is the list of StatefulSets

表 6-218 Job 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	Yes	ObjectMeta object	Standard list metadata.
spec	Yes	JobSpec object	Specification of the desired behavior of a job
status	No	JobStatus object	Current status of a job

表 6-219 JobStatus 字段数据结构说明

参数	是否必选	参数类型	描述
active	No	Integer	The number of actively running pods.
completionTime	No	Time	Replicas is the number of actual replicas.
conditions	No	Array of JobCondition objects	The latest available observations of an object's current state. More info:
failed	No	Integer	The number of pods which reached phase Failed.
startTime	No	Time	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
succeeded	No	Integer	The number of pods which reached phase Succeeded.

表 6-220 JobSpec 字段数据结构说明

参数	是否必选	参数类型	描述
activeDeadlineSeconds	No	Integer	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer
backoffLimit	No	Integer	Specifies the number of retries before marking this job failed. Defaults to 6.
priority	No	Integer	任务优先级，数值越大优先级越高，默认值为0。 取值范围：[-10, 10]
completions	No	Integer	Specifies the desired number of successfully finished pods the job should be run with. Setting to nil means that the success of any pod signals the success of all pods, and allows parallelism to have any positive value. Setting to 1 means that parallelism is limited to 1 and the success of that pod signals the success of the job.
manualSelector	No	Boolean	manualSelector controls generation of pod labels and pod selectors. Leave manualSelector unset unless you are certain what you are doing. When false or unset, the system picks labels unique to this job and appends those labels to the pod template. When true, the user is responsible for picking unique labels and specifying the selector. Failure to pick a unique label may cause this and other jobs to not function correctly. However, You may see manualSelector=true in jobs that were created with the old extensions/v1beta1 API.

参数	是否必选	参数类型	描述
parallelism	No	Integer	Specifies the maximum desired number of pods the job should run at any given time. The actual number of pods running in steady state will be less than this number when $(.spec.completions - .status.successful) < .spec.parallelism$, i.e. when the work left to do is less than max parallelism.
selector	Yes	labelSelector object	Selector is a label query over pods that should match the replica count. If empty, defaulted to labels on the pod template.
template	Yes	PodTemplateSpec object	Template is the object that describes the pod that will be created if insufficient replicas are detected. Each pod stamped out by the StatefulSet will fulfill this Template, but have a unique identity from the rest of the StatefulSet.

表 6-221 JobCondition 字段数据结构说明

参数	是否必选	参数类型	描述
lastProbeTime	No	String	Last time the condition was checked.
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.
message	No	String	Human readable message indicating details about last transition
reason	No	String	(brief) reason for the condition's last transition.
status	No	String	Status of the condition, one of True, False, Unknown.
type	No	String	Type of job condition, Complete or Failed.

表 6-222 Service v1 core 数据结构描述

参数	是否必选	参数类型	描述
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Service .
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	Yes	ObjectMeta object	-
spec	Yes	ServiceSpec object	-
status	No	ServiceStatus object	-

表 6-223 ServiceSpec 字段数据结构说明

参数	是否必选	参数类型	描述
ports	Yes	Array of ServicePort objects	The list of ports that are exposed by this service.
selector	No	Object	This service will route traffic to pods having labels matching this selector. Label keys and values that must match in order to receive traffic for this service. If empty, all pods are selected, if not specified, endpoints must be manually specified.

参数	是否必选	参数类型	描述
clusterIP	No	String	<p>clusterIP is the IP address of the service and is usually assigned randomly by the master. If an address is specified manually and is not in use by others, it will be allocated to the service; otherwise, creation of the service will fail. This field cannot be changed through updates. Valid values are "None", empty string (""), or a valid IP address.</p> <p>"None" can be specified for headless services when proxying is not required. Only applies to types ClusterIP, NodePort, and LoadBalancer. Ignored if type is ExternalName.</p>
type	No	String	<p>type determines how the Service is exposed. Defaults to ClusterIP. Valid options are ExternalName, ClusterIP and LoadBalancer.</p> <p>"ExternalName" maps to the specified externalName.</p> <p>"ClusterIP" allocates a cluster-internal IP address for load-balancing to endpoints. Endpoints are determined by the selector or if that is not specified, by manual construction of an Endpoints object. If clusterIP is "None", no virtual IP is allocated and the endpoints are published as a set of endpoints rather than a stable IP.</p> <p>"LoadBalancer" builds on NodePort and creates an external load-balancer (if supported in the current cloud) which routes to the clusterIP.</p> <p>说明 社区版本支持nodePort类型的service，CCI场景下不支持该类型的service</p>

参数	是否必选	参数类型	描述
externalIPs	No	Array of strings	externalIPs is a list of IP addresses for which nodes in the cluster will also accept traffic for this service. These IPs are not managed by Kubernetes. The user is responsible for ensuring that traffic arrives at a node with this IP. A common example is external load-balancers that are not part of the Kubernetes system.
externalTrafficPolicy	No	String	externalTrafficPolicy denotes if this Service desires to route external traffic to node-local or cluster-wide endpoints. valid values are "Local" and "Cluster" <ul style="list-style-type: none">- "Local" preserves the client source IP and avoids a second hop for LoadBalancer and Nodeport type services, but risks potentially imbalanced traffic spreading.- "Cluster" obscures the client source IP and may cause a second hop to another node, but should have good overall load-spreading.
healthCheckNodePort	No	Integer	healthCheckNodePort specifies the healthcheck nodePort for the service. If not specified, HealthCheckNodePort is created by the service api backend with the allocated nodePort. Will use user-specified nodePort value if specified by the client. Only effects when Type is set to LoadBalancer and ExternalTrafficPolicy is set to Local.

参数	是否必选	参数类型	描述
externalName	No	String	externalName is the external reference that kubedns or equivalent will return as a CNAME record for this service. No proxying will be involved. Must be a valid DNS name and requires Type to be ExternalName.
sessionAffinity	No	String	Used to maintain session affinity. Enable client IP based session affinity. Must be ClientIP or None. Defaults to None.
loadBalancerIP	No	String	Only applies to Service Type: LoadBalancer LoadBalancer will get created with the IP specified in this field. This feature depends on whether the underlying cloud-provider supports specifying the loadBalancerIP when a load balancer is created. This field will be ignored if the cloud-provider does not support the feature.
loadBalancerSourceRanges	No	Array of strings	Optional: If specified and supported by the platform, this will restrict traffic through the cloud-provider.load-balancer will be restricted to the specified client IPs. This field will be ignored if the cloud-provider does not support the feature.

参数	是否必选	参数类型	描述
publishNotReadyAddresses	No	Boolean	publishNotReadyAddresses, when set to true, indicates that DNS implementations must publish the notReadyAddresses of subsets for the Endpoints associated with the Service. The default value is false. The primary use case for setting this field is to use a StatefulSet's Headless Service to propagate SRV records for its Pods without respect to their readiness for purpose of peer discovery. This field will replace the service.alpha.kubernetes.io/tolerate-unready-endpoints when that annotation is deprecated and all clients have been converted to use this field.
sessionAffinityConfig	No	SessionAffinityConfig object	sessionAffinityConfig contains the configurations of session affinity.

表 6-224 ServiceStatus 字段数据结构说明

参数	是否必选	参数类型	描述
loadBalancer	No	loadBalancerStatus object	LoadBalancer contains the current status of the load-balancer, if one is present

表 6-225 ServicePort 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	<p>The name of this port within the service. This must be a DNS_LABEL. All ports within a ServiceSpec must have unique names. This maps to the 'Name' field in EndpointPort objects. Optional if only one ServicePort is defined on this service.</p> <p>Value length: 0 character < String length ≤ 63 characters.</p> <p>The string must comply with regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.</p>
protocol	No	String	<p>The IP protocol for this port. Supports "TCP" and "UDP".</p> <p>This parameter can be set to:</p> <ul style="list-style-type: none">• TCP• UDP
port	Yes	Integer	<p>The port that will be exposed by this service.</p> <p>Value range: (0,65535].</p>
targetPort	No	String	<p>Number or name of the port to access on the pods targeted by the service. Number must be in the range 1 to 65535. Name must be an IANA_SVC_NAME. If this is a string, it will be looked up as a named port in the target Pod's container ports. If this is not specified, the value of Port is used (an identity map). Defaults to the service port.</p> <p>Value range: (0,65535].</p>

参数	是否必选	参数类型	描述
nodePort	No	Integer	The port on each node on which this service is exposed when type=NodePort or LoadBalancer. Usually assigned by the system. If specified, it will be allocated to the service if unused or else creation of the service will fail. Default is to auto-allocate a port if the ServiceType of this Service requires one. Value range: [30000,32767].

表 6-226 loadBalancerStatus 数据结构说明

参数	是否必选	参数类型	描述
ingress	No	Array of LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 6-227 LoadBalancerIngress 字段数据结构说明

参数	是否必选	参数类型	描述
ip	No	String	IP is set for load-balancer ingress points that are IP based.
hostname	No	String	Hostname is set for load-balancer ingress points that are DNS based.

表 6-228 SessionAffinityConfig 字段数据结构说明

参数	是否必选	参数类型	描述
clientIP	No	ClientIPConfig object	clientIP contains the configurations of Client IP based session affinity.

表 6-229 ClientIPConfig 字段数据结构说明

参数	是否必选	参数类型	描述
timeoutSeconds	No	Integer	timeoutSeconds specifies the seconds of ClientIP type session sticky time. The value must be >0 && <=86400(for 1 day) if ServiceAffinity == "ClientIP". Default value is 10800(for 3 hours).

表 6-230 ServiceList 数据结构说明

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	ListMeta object	Standard list metadata.
items	Array of Service objects	List of services.

表 6-231 Ingress v1beta1 extensions 数据结构

参数	是否必选	参数类型	描述
apiVersion	No	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.

参数	是否必选	参数类型	描述
kind	No	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	No	ObjectMeta object	Standard object's metadata.
spec	No	IngressSpec object	Spec is the desired state of the Ingress.
status	No	IngressStatus	Status is the current state of the Ingress.

表 6-232 IngressSpec 字段数据结构说明

参数	是否必选	参数类型	描述
backend	No	IngressBackend object	A default backend capable of servicing requests that do not match any rule. At least one of <i>backend</i> or <i>rules</i> must be specified. This field is optional to allow the loadbalancer controller or defaulting logic to specify a global default.
rules	No	Array of IngressRule objects	A list of host rules used to configure the Ingress. If unspecified, or no rule matches, all traffic is sent to the default backend.
tls	No	Array of IngressTLS objects	TLS configuration. Currently the Ingress only supports a single TLS port, 443. If multiple members of this list specify different hosts, they will be multiplexed on the same port according to the hostname specified through the SNI TLS extension, if the ingress controller fulfilling the ingress supports SNI.

表 6-233 IngressStatus 字段数据结构说明

参数	是否必选	参数类型	描述
loadBalancer	No	loadBalancer Status object	LoadBalancer contains the current status of the load-balancer.

表 6-234 IngressBackend 字段数据结构说明

参数	是否必选	参数类型	描述
serviceName	No	String	Specifies the name of the referenced service.
servicePort	No	String	Specifies the port of the referenced service.

表 6-235 IngressTLS 数据结构说明

参数	是否必选	参数类型	描述
hosts	No	Array of strings	Hosts are a list of hosts included in the TLS certificate. The values in this list must match the name/s used in the tlsSecret. Defaults to the wildcard host setting for the loadbalancer controller fulfilling this Ingress, if left unspecified.
secretName	No	String	SecretName is the name of the secret used to terminate SSL traffic on 443. Field is left optional to allow SSL routing based on SNI hostname alone. If the SNI host in a listener conflicts with the "Host" header field used by an IngressRule, the SNI host is used for termination and value of the Host header is used for routing

表 6-236 IngressRule 数据结构说明

参数	是否必选	参数类型	描述
host	No	String	ost is the fully qualified domain name of a network host, as defined by RFC 3986. Note the following deviations from the "host" part of the URI as defined in the RFC: 1. IPs are not allowed. Currently an IngressRuleValue can only apply to the IP in the Spec of the parent Ingress. 2. The : delimiter is not respected because ports are not allowed. Currently the port of an Ingress is implicitly :80 for http and :443 for https. Both these may change in the future. Incoming requests are matched against the host before the IngressRuleValue. If the host is unspecified, the Ingress routes all traffic based on the specified IngressRuleValue
http	No	IngressRuleValue object	IngressRuleValue represents a rule to route requests for this IngressRule.If unspecified, the rule defaults to an http catch-all. Whether that sends just traffic matching the host to the default backend or all traffic to the default backend, is left to the controller fulfilling the Ingress. Http is currently the only supported IngressRuleValue.

表 6-237 IngressRuleValue 数据结构说明

参数	是否必选	参数类型	描述
http	No	HTTPIngressRuleValue object	HTTP ingress rule.

表 6-238 HTTPIngressRuleValue 数据结构说明

参数	是否必选	参数类型	描述
paths	No	Array of HTTPIngressPath objects	A collection of paths that map requests to backends.

表 6-239 HTTPIngressPath 数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Path is an extended POSIX regex as defined by IEEE Std 1003.1, (i.e this follows the egrep/unix syntax, not the perl syntax) matched against the path of an incoming request. Currently it can contain characters disallowed from the conventional "path" part of a URL as defined by RFC 3986. Paths must begin with a '/'. If unspecified, the path defaults to a catch all sending traffic to the backend
backend	Yes	Ingress Backend object	Backend defines the referenced service endpoint to which the traffic will be forwarded to
property	Yes	Object	Extension property on the path

表 6-240 loadBalancerStatus 字段数据结构说明

参数	是否必选	参数类型	描述
ingress	No	Array of LoadBalancerIngress objects	Ingress is a list containing ingress points for the load-balancer. Traffic intended for the service should be sent to these ingress points.

表 6-241 IngressList 数据结构说明

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.

参数	参数类型	描述
metadata	ListMeta object	Standard list metadata.
items	Array of Ingress v1beta1 extensions objects	List of Ingress.

表 6-242 Configmap v1 core 请求参数

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	Yes	ObjectMeta object	Standard list metadata.
data	Yes	Object	Data contains the configuration data. Each key must consist of alphanumeric characters, '-', '_' or '.'. value值最大长度为512.

表 6-243 ConfigmapList 数据结构说明

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	ListMeta object	Standard list metadata.
items	Array of Configmap v1 core objects	List of Configmap.

表 6-244 Secret v1 core 数据结构说明

参数	是否必选	参数类型	描述
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Secret .
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	Yes	ObjectMeta object	-
data	No	Object	Data contains the secret data. Each key must consist of alphanumeric characters, '-', '_' or '!'. The serialized form of the secret data is a base64 encoded string , representing the arbitrary (possibly non-string) data value here
stringData	No	Object	stringData allows specifying non-binary secret data in string form. It is provided as a write-only convenience method. All keys and values are merged into the data field on write, overwriting any existing values. It is never output when reading from the API

参数	是否必选	参数类型	描述
type	No	String	Used to facilitate programmatic handling of secret data. The primitive k8s supports the following secret types, for details, see 表6-245 . <ul style="list-style-type: none"> • Opaque • kubernetes.io/dockercfg • kubernetes.io/dockerconfigjson • kubernetes.io/tls

表 6-245 不同类型 secret 对应的 data 中 key 值约束

secret类型	data中必须含有的key值	描述
Opaque	无	secret type Opaque is the default; arbitrary user-defined data
kubernetes.io/dockercfg	.dockercfg	secret type kubernetes.io/dockercfg contains a dockercfg file that follows the same format rules as ~/.dockercfg
kubernetes.io/tls	tls.key tls.crt	secret type kubernetes.io/tls contains information about a TLS client or server secret. It is primarily used with TLS termination of the Ingress resource, but may be used in other types.
kubernetes.io/dockerconfigjson	.dockerconfigjson	SecretTypeDockerConfigJson contains a dockercfg file that follows the same format rules as ~/.docker/config.json

表 6-246 ServiceList 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	ListMeta object	Standard list metadata.
items	Array of Secret v1 core objects	List of Secrets.

表 6-247 PersistentVolumeClaimList v1 core 数据结构说明

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
metadata	ListMeta object	Standard list metadata.
items	Array of PersistentVolumeClaim object	A list of persistent volume claims.

表 6-248 Event v1 core 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.

参数	参数类型	描述
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
count	integer	The number of times this event has occurred.
firstTimestamp	Time	The time at which the event was first recorded. (Time of server receipt is in TypeMeta.)
involvedObject	involvedObject object	The object that this event is about.
lastTimestamp	Time	The time at which the most recent occurrence of this event was recorded.
message	String	A human-readable description of the status of this operation.
metadata	metadata object	Standard object's metadata.
reason	String	This should be a short, machine understandable string that gives the reason for the transition into the object's current status.
source	EventSource object	The component reporting this event. Should be a short machine understandable string.
type	String	Type of this event (Normal, Warning), new types could be added in the future
eventTime	time.Time	Time when this Event was first observed.
Series	EventSeries object	Data about the Event series this event represents or nil if it is a singleton Event.
action	String	What action was taken/failed regarding to the Regarding object.
related	ObjectReference object	Optional secondary object for more complex actions.
reportingComponent	String	Name of the controller that emitted this Event, e.g. `kubernetes.io/kubelet`
reportingInstance	String	ID of the controller instance, e.g. `kubelet-xyzf`.

表 6-249 involvedObject 字段数据结构说明

参数	参数类型	描述
kind	String	Kind of the referent.
namespace	String	Namespace of the referent.
name	String	Name of the referent.
uid	String	UID of the referent.
apiVersion	String	API version of the referent.
resourceVersion	String	Specific resourceVersion to which this reference is made.
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.

表 6-250 EventSource 字段数据结构说明

参数	参数类型	描述
component	String	Component from which the event is generated.
host	String	Node name on which the event is generated.

表 6-251 EventSerie 数据结构说明

参数	参数类型	描述
count	<i>integer</i>	Number of occurrences in this series up to the last heartbeat time

参数	参数类型	描述
lastObservedTime	time.Time	Time of the last occurrence observed
state	String	State of this Series: Ongoing 、 Finished or Unknown

表 6-252 ObjectReference 字段数据结构说明

参数	参数类型	描述
apiVersion	String	API version of the referent.
fieldPath	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as desiredState.manifest.containers[2]. For example, if the object reference is to a container within a pod, this would take on a value like: "spec.containers{name}" (where "name" refers to the name of the container that triggered the event) or if no container name is specified "spec.containers[2]" (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	String	Kind of the referent.
name	String	Name of the referent.
namespace	String	Namespace of the referent.
resourceVersion	String	Specific resourceVersion to which this reference is made, if any.
uid	String	UID of the referent.

表 6-253 EventList v1 core 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.

参数	参数类型	描述
items	Array of Event v1 core objects	List of services
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	ListMeta object	Standard list metadata.

表 6-254 LocalDirVolumeSource 数据结构说明

参数	参数类型	描述
sizeLimit	Object	所使用的localDir存储大小。

表 6-255 emptyDir 数据结构说明

参数	参数类型	描述
sizeLimit	integer	所使用的emptyDir存储大小。 取值：(0, 2147483647] 单位：Gi。
medium	String	介质类型，可设置为： <ul style="list-style-type: none">LocalVolume：使用磁盘类型为超高IO的云硬盘。LocalSSD：使用本地SSD。 说明 若不设置，默认使用超高IO的云硬盘。

表 6-256 Endpoints 数据结构说明

参数	参数类型	描述
kind	String	String Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Secret .
apiVersion	String	String APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	ObjectMetadata object	-
subsets	Array of EndpointSubset objects	The set of all endpoints is the union of all subsets. Addresses are placed into subsets according to the IPs they share. A single address with multiple ports, some of which are ready and some of which are not (because they come from different containers) will result in the address being displayed in different subsets for the different ports. No address will appear in both Addresses and NotReadyAddresses in the same subset. Sets of addresses and ports that comprise a service.

表 6-257 EndpointSubset 数据结构说明

参数	参数类型	描述
addresses	Array of EndpointAddress objects	IP addresses which offer the related ports that are marked as ready. These endpoints should be considered safe for load balancers and clients to utilize.

参数	参数类型	描述
notReadyAddresses	Array of EndpointAddress objects	IP addresses which offer the related ports but are not currently marked as ready because they have not yet finished starting, have recently failed a readiness check, or have recently failed a liveness check.
ports	Array of EndpointPort objects	Port numbers available on the related IP addresses.

表 6-258 EndpointAddress 数据结构说明

参数	参数类型	描述
ip	String	The IP of this endpoint. May not be loopback (127.0.0.0/8), link-local (169.254.0.0/16), or link-local multicast ((224.0.0.0/24) IPv6 is also accepted but not fully supported on all platforms. Also, certain kubernetes components, like kube-proxy, are not IPv6 ready
hostname	String	The Hostname of this endpoint
nodename	String	Optional: Node hosting this endpoint. This can be used to determine endpoints local to a node.
targetRef	ObjectReference object	Reference to object providing the endpoint.
nodeAvailabilityZone	String	Optional: The availability zone of the endpoint's host node

表 6-259 EndpointPort 数据结构说明

参数	参数类型	描述
name	String	The name of this port (corresponds to ServicePort.Name). Must be a DNS_LABEL. Optional only if one port is defined.
port	Integer	The port number of the endpoint.
protocol	String	The IP protocol for this port. Must be UDP or TCP. Default is TCP

表 6-260 EndpointsList v1 core 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
items	Array of Endpoints objects	List of endpoints
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	ListMeta object	Standard list metadata.

表 6-261 ReplicaSet 数据结构说明

参数	参数类型	描述
kind	String	String Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is Secret .

参数	参数类型	描述
apiVersion	String	String APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
metadata	ObjectMeta object	-
spec	ReplicaSetSpec object	Spec defines the specification of the desired behavior of the ReplicaSet.
status	ReplicaSetStatus object	Status is the most recently observed status of the ReplicaSet. This data may be out of date by some window of time. Populated by the system. Read-only.

表 6-262 ReplicaSetSpec 数据结构说明

参数	参数类型	描述
replicas	Integer	Replicas is the number of desired replicas. This is a pointer to distinguish between explicit zero and unspecified. Defaults to 1.
minReadySeconds	Integer	Minimum number of seconds for which a newly created pod should be ready without any of its container crashing, for it to be considered available。 Defaults to 0 (pod will be considered available as soon as it is ready)
selector	LabelSelector object	Selector is a label query over pods that should match the replica count. Label keys and values that must match in order to be controlled by this replica set. It must match the pod template's labels.
template	PodTemplateSpec object	Template is the object that describes the pod that will be created if insufficient replicas are detected.

表 6-263 ReplicaSetStatus 数据结构说明

参数	参数类型	描述
replicas	Integer	Replicas is the most recently observed number of replicas.
fullyLabeledReplicas	Integer	The number of pods that have labels matching the labels of the pod template of the replicaset.
readyReplicas	Integer	The number of ready replicas for this replica set.
availableReplicas	Integer	The number of available replicas (ready for at least minReadySeconds) for this replica set.
observedGeneration	Integer	ObservedGeneration reflects the generation of the most recently observed ReplicaSet.
conditions	ReplicaSetCondition object	Represents the latest available observations of a replica set's current state.

表 6-264 ReplicaSetCondition 数据结构说明

参数	参数类型	描述
type	String	Type of replica set condition. Available Values: -- ReplicaFailure: ReplicaSetReplicaFailure is added in a replica set when one of its pods fails to be created due to insufficient quota, limit ranges, pod security policy, node selectors, etc. or deleted due to kubelet being down or finalizers are failing.
status	String	Status of the condition, one of True, False, Unknown.
lastTransitionTime	Object	The last time the condition transitioned from one status to another.
reason	String	The reason for the condition's last transition.
message	String	A human readable message indicating details about the transition.

表 6-265 ReplicaSetList v1 core 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
items	Array of ReplicaSet objects	List of endpoints
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	ListMeta object	Standard list metadata.

表 6-266 Volcano Job batch_v1alpha1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#resources .
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#types-kinds .
metadata	Yes	ObjectMeta object	Standard object's metadata. More info: https://git.k8s.io/community/contributors/devel/api-conventions.md#metadata .

参数	是否必选	参数类型	描述
spec	Yes	VolcanoJobSpec object	Specification of the desired behavior of a cron job, including the minAvailable.
status	No	VolcanoJobStatus object	Current status of Job.

表 6-267 VolcanoJobSpec 字段数据结构说明

参数	是否必选	参数类型	描述
maxRetry	No	Integer	The limit for retrying submitting job, default is 3.
minAvailable	Yes	Integer	The minimal available pods to run for this Job. 该参数值大于0。
plugins	No	VolcanoPlugin object	Enabled task plugins when creating job.
policies	No	Array of VolcanoJobPolicy objects	Specifies the default lifecycle of tasks.
queue	No	String	The name of the queue on which job should be created.
schedulerName	No	String	SchedulerName is the default value of `tasks.template.spec.schedulerName`.
tasks	Yes	Array of VolcanoJobTask objects	Tasks specifies the task specification of Job.
volumes	No	Array of VolcanoJobVolume objects	The volumes for Job.

表 6-268 VolcanoJobStatus 字段数据结构说明

参数	是否必选	参数类型	描述
ControlledResources	Yes	Object	All of the resources that are controlled by this job. { "plugin-env": "env", "plugin-ssh": "ssh", "plugin-svc": "svc" }.

参数	是否必选	参数类型	描述
Succeeded	No	Integer	The number of pods which reached phase Succeeded.
failed	No	Integer	The number of pods which reached phase Failed.
minAvailable	Yes	Integer	The minimal available pods to run for this Job.
pending	No	Integer	The number of pending pods.
retryCount	No	Integer	The number that volcano retried to submit the job.
running	No	Integer	The number of running pods.
version	No	Integer	Job's current version.
state	Yes	VolcanoJobStatusState object	Current state of Job.

表 6-269 VolcanoJobPolicy 字段数据结构说明

参数	是否必选	参数类型	描述
action	Yes	String	The action that will be taken to the PodGroup according to Event. One of \"Restart\", \"None\". Default to None.
event	Yes	String	The Event recorded by scheduler; the controller takes actions according to this Event.
timeout	No	Object	Timeout is the grace period for controller to take actions. Default to nil (take action immediately).

表 6-270 VolcanoJobTask 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name specifies the name of tasks.
policies	No	VolcanoJobPolicy object	Specifies the lifecycle of task.
replicas	No	Integer	Replicas specifies the replicas of this TaskSpec in Job.

表 6-273 VolcanoTaskVolumeClaimSpec 字段数据结构说明

参数	是否必选	参数类型	描述
accessModes	Yes	Array of strings	AccessModes contains the desired access modes the volume should have. ReadWriteOnce – the volume can be mounted as read-write by a single node ReadOnlyMany – the volume can be mounted read-only by many nodes ReadWriteMany – the volume can be mounted as read-write by many nodes
resources	Yes	ResourceRequirements object	Resources represents the minimum resources the volume should have.
storageClassName	Yes	String	Name of the StorageClass required by the claim. 支持如下字段： <ul style="list-style-type: none"> 云硬盘存储卷（EVS）当前支持高I/O（sas）、超高I/O（ssd）和普通I/O（sata）。 文件存储卷（SFS）当前支持标准文件协议类型（nfs-rw）。

表 6-274 VolcanoPlugin 字段数据结构说明

参数	是否必选	参数类型	描述
env	No	Array of strings	Set VK_TASK_INDEX to each container, is a index for giving the identity to container.
svc	No	Array of strings	Create Service and *.host to enable pods communicate.
ssh	No	Array of strings	Sign in ssh without password, e.g. use command mpirun or mpiexec. 可填写值“no-root”表明ssh登录为非root用户登录。

表 6-275 TFJob kubeflow_v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is TFJob .
metadata	Yes	ObjectMeta object	Standard object's metadata.
spec	Yes	TFJobSpec object	Specification of the desired behavior of a tfjob.
status	No	JobStatus object	Current status of TFJob.

表 6-276 TFJobSpec 字段数据结构说明

参数	是否必选	参数类型	描述
activeDeadlineSeconds	No	Integer	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer.
backoffLimit	No	Integer	Optional number of retries before marking this job failed.
cleanPodPolicy	No	CleanPodPolicy object	CleanPodPolicy defines the policy to kill pods after TFJob is succeeded. Default to Running.
ttlSecondsAfterFinished	No	Integer	TTLSecondsAfterFinished is the TTL to clean up tf-jobs (temporary before kubernetes adds the cleanup controller). It may take extra ReconcilePeriod seconds for the cleanup, since reconcile gets called periodically.

参数	是否必选	参数类型	描述
tfReplicaSpecs	Yes	Array of ReplicaSpec objects	TFReplicaSpecs is map of TFReplicaType and ReplicaSpec specifies the TF replicas to run. For example, { "PS": ReplicaSpec, "Worker": ReplicaSpec, }

表 6-277 CleanPodPolicy 字段可选值说明

可选值	描述
All	When the job is finished, kill all pods that the job created.
Running	When the job is finished, kill pods that the job created and is in running phase.
None	When the job is finished, do not kill any pods that the job created.

表 6-278 TFReplicaType 字段可选值说明

可选值	描述
PS	PS is the type for parameter servers of distributed TensorFlow.
Worker	Worker is the type for workers of distributed TensorFlow. This is also used for non-distributed TensorFlow.
Chief	Chief is the type for chief worker of distributed TensorFlow. If there is "chief" replica type, it's the "chief worker". Else, worker:0 is the chief worker.
Evaluator	Evaluator is the type for evaluation replica in TensorFlow.

表 6-279 ReplicaSpec 字段数据结构说明

参数	是否必选	参数类型	描述
replicas	No	Integer	Replicas is the desired number of replicas of the given template. If unspecified, defaults to 1.
template	Yes	PodTemplateSpec object	Template is the object that describes the pod that will be created for this replica. RestartPolicy in PodTemplateSpec will be override by RestartPolicy in ReplicaSpec.
restartPolicy	No	String	Restart policy for all replicas within the job. One of Always, OnFailure, Never and ExitCode. Default to Never.

表 6-280 JobStatus 字段数据结构说明

参数	是否必选	参数类型	描述
conditions	No	Array of JobCondition objects	Conditions is an array of current observed job conditions.
replicaStatuses	No	Array of ReplicaStatus objects	ReplicaStatuses is map of ReplicaType and ReplicaStatus, specifies the status of each replica.
startTime	No	Time	Represents time when the job was acknowledged by the job controller. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
completionTime	No	Time	Represents time when the job was completed. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.
lastReconcileTime	No	Time	Represents last time when the job was reconciled. It is not guaranteed to be set in happens-before order across separate operations. It is represented in RFC3339 form and is in UTC.

表 6-281 JobCondition 字段数据结构说明

参数	是否必选	参数类型	描述
type	No	String	Type of job condition. 说明 Job的状态有: -Created: means the job has been accepted by the system, but one or more of the pods/services has not been started. This includes time before pods being scheduled and launched. -Running: means all sub-resources (e.g. services/pods) of this job have been successfully scheduled and launched. The training is running without error. -Restarting: means one or more sub-resources (e.g. services/pods) of this job reached phase failed but maybe restarted according to it's restart policy which specified by user in v1.PodTemplateSpec. The training is freezing/pending. -Succeeded: means all sub-resources (e.g. services/pods) of this job reached phase have terminated in success. The training is complete without error. -Failed: means one or more sub-resources (e.g. services/pods) of this job reached phase failed with no restarting. The training has failed its execution.
status	No	String	Status of the condition, one of True, False, Unknown.
reason	No	String	(brief) reason for the condition's last transition.
message	No	String	Human readable message indicating details about last transition
lastUpdateTime	No	Time	The last time this condition was updated.
lastTransitionTime	No	Time	Last time the condition transitioned from one status to another.

表 6-282 ReplicaStatus 字段数据结构说明

参数	是否必选	参数类型	描述
active	No	Integer	The number of actively running pods.

参数	是否必选	参数类型	描述
succeeded	No	Integer	The number of pods which reached phase Succeeded.
failed	No	Integer	The number of pods which reached phase Failed.

表 6-283 MXJob kubeflow_v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is MXJob .
metadata	Yes	ObjectMetadata object	Standard object's metadata.
spec	Yes	MXJobSpec	Specification of the desired behavior of a mxjob.
status	No	JobStatus object	Current status of MXJob.

表 6-284 MXJobSpec 字段数据结构说明

参数	是否必选	参数类型	描述
activeDeadlineSeconds	No	Integer	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer.
backoffLimit	No	Integer	Optional number of retries before marking this job failed.

参数	是否必选	参数类型	描述
cleanPodPolicy	No	CleanPodPolicy object	CleanPodPolicy defines the policy to kill pods after TFJob is succeeded. Default to Running.
ttlSecondsAfterFinished	No	Integer	TTLSecondsAfterFinished is the TTL to clean up tf-jobs (temporary before kubernetes adds the cleanup controller). It may take extra ReconcilePeriod seconds for the cleanup, since reconcile gets called periodically.
mxReplicaSpecs	Yes	Array of ReplicaSpec objects	MXReplicaSpecs is map of MXReplicaType and MXReplicaSpec specifies the MX replicas to run. For example, { "Scheduler": MXReplicaSpec, "Server": MXReplicaSpec, "Worker": MXReplicaSpec, }

表 6-285 MXReplicaType 字段可选值说明

可选值	描述
Scheduler	Scheduler is the type for scheduler replica in MXNet.
Worker	Worker is the type for workers of distributed MXNet.
Server	Server is the type for parameter servers of distributed MXNet.

表 6-286 PyTorchJob kubeflow_v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1 .
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is PyTorchJob .
metadata	Yes	ObjectMeta object	Standard object's metadata.
spec	Yes	PyTorchJobSpec object	Specification of the desired behavior of a pytorchjob.
status	No	JobStatus object	Current status of pytorchJob.

表 6-287 PyTorchJobSpec 字段数据结构说明

参数	是否必选	参数类型	描述
activeDeadlineSeconds	No	Integer	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer.
backoffLimit	No	Integer	Optional number of retries before marking this job failed.
cleanPodPolicy	No	CleanPodPolicy object	CleanPodPolicy defines the policy to kill pods after TFJob is succeeded. Default to Running.
ttlSecondsAfterFinished	No	Integer	TTLSecondsAfterFinished is the TTL to clean up tf-jobs (temporary before kubernetes adds the cleanup controller). It may take extra ReconcilePeriod seconds for the cleanup, since reconcile gets called periodically.

参数	是否必选	参数类型	描述
ReplicaSpecs	Yes	Array of ReplicaSpec objects	A map of PyTorchReplicaType (type) to ReplicaSpec (value). Specifies the PyTorch cluster configuration. For example, { "Master": PyTorchReplicaSpec, "Worker": PyTorchReplicaSpec, }

表 6-288 PytorchReplicaType 字段可选值说明

可选值	描述
Master	Master is the type of Master of distributed PyTorch.
Worker	Worker is the type for workers of distributed PyTorch.

表 6-289 TFJobList kubeflow_v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	No	ListMeta object	Standard type metadata.
items	Yes	Array of TFJob kubeflow_v1 objects	Lists of tfjobs.

表 6-290 MXJobList kubeflow_v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	No	ListMeta object	Standard type metadata.
items	Yes	Array of MXJob kubeflow_v1 objects	Lists of mxjobs.

表 6-291 PyTorchJobList kubeflow_v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	No	ListMeta object	Standard type metadata.
items	Yes	Array of PyTorchJob kubeflow_v1 objects	Lists of pytorchjobs.

表 6-292 MPIJob kubeflow_v1alpha2 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values. The value of this parameter is v1alpha2 .
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase. The value of this parameter is MPIJob .
metadata	Yes	ObjectMeta object	Standard object's metadata.
spec	Yes	MPIJobSpec object	Specification of the desired behavior of a mpijob.
status	No	JobStatus object	Current status of MPIJob.

表 6-293 MPIJobSpec 字段数据结构说明

参数	是否必选	参数类型	描述
activeDeadlineSecond	No	Integer	Specifies the duration in seconds relative to the startTime that the job may be active before the system tries to terminate it; value must be positive integer.
backoffLimit	No	Integer	Optional number of retries before marking this job failed.
cleanPodPolicy	No	CleanPodPolicy object	CleanPodPolicy defines the policy to kill pods after TFJob is succeeded. Default to Running.
slotsPerWorker	No	Integer	Specifies the number of slots per worker used in hostfile. Defaults to 1.

参数	是否必选	参数类型	描述
mpiReplicaSpecs	Yes	Array of ReplicaSpec objects	MPIReplicaSpecs is map of MPIReplicaType and MPIReplicaSpec specifies the MPI replicas to run. For example, { "Launcher": MPIReplicaSpec, "Worker": MPIReplicaSpec, }

表 6-294 MPIReplicaType 字段可选值说明

可选值	描述
Launcher	Launcher is the type for launcher replica in MPI.
Worker	Worker is the type for workers of distributed MPI.

表 6-295 MPIJobList kubeflow_v1alpha2 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	No	ListMeta object	Standard type metadata.
items	Yes	Array of MPIJob kubeflow_v1alpha2 objects	Lists of mpijobs.

表 6-296 PersistentVolumeClaim v1 数据结构说明

参数	参数类型	描述
apiVersion	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	ListMeta v1 meta object	Standard object's metadata.
spec	PersistentVolumeClaimSpec object	Spec defines the desired characteristics of a volume requeste PersistentVolumeClaimStatus d by a pod author.
status	PersistentVolumeClaimStatus object	Status represents the current information/status of a persistent volume claim. Read-only.

表 6-297 ListMeta v1 meta 数据结构说明

名称	参数类型	描述
continue	string	continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response
resourceVersion	string	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only
selfLink	string	SelfLink is a URL representing this object. Populated by the system. Read-only

表 6-298 PersistentVolumeClaimSpec 字段数据结构说明

参数	是否必选	参数类型	描述
volumeName	No	String	VolumeName is the binding reference to the PersistentVolume backing this claim.
accessModes	Yes	Array of strings	AccessModes contains the desired access modes the volume should have. ReadWriteOnce – the volume can be mounted as read-write by a single node ReadOnlyMany – the volume can be mounted read-only by many nodes ReadWriteMany – the volume can be mounted as read-write by many nodes
resources	Yes	ResourceRequirements object	Resources represents the minimum resources the volume should have.
selector	No	labelSelector object	A label query over volumes to consider for binding.
storageClassName	Yes	String	Name of the StorageClass required by the claim. 支持如下字段： <ul style="list-style-type: none"> 云硬盘存储卷（EVS） 当前支持高I/O（sas）、超高I/O（ssd）和普通I/O（sata）。 文件存储卷（SFS） 当前支持标准文件协议类型（nfs-rw）。 极速文件存储卷（SFS Turbo） 当前支持性能型极速文件存储卷（efs-performance）和标准型极速文件存储卷（efs-standard）。
volumeMode	No	String	volumeMode defines what type of volume is required by the claim. Can be: <ul style="list-style-type: none"> - Block: the volume will not be formatted with a filesystem and will remain a raw block device - Filesystem: the volume will be or is formatted with a filesystem

表 6-299 PersistentVolumeClaimStatus 字段数据结构说明

参数	是否必选	参数类型	描述
accessModes	No	Array of strings	AccessModes contains the actual access modes the volume backing the PVC has. ReadWriteOnce - can be mounted read/write mode to exactly 1 host ReadOnlyMany - can be mounted in read-only mode to many hosts ReadWriteMany - can be mounted in read/write mode to many hosts
capacity	No	Array of ResourceRequirements objects	Represents the actual resources of the underlying volume.
phase	No	String	Phase represents the current phase of PersistentVolumeClaim. pending - used for PersistentVolumeClaims that are not yet bound Bound - used for PersistentVolumeClaims that are bound Lost - used for PersistentVolumeClaims that lost their underlying PersistentVolume. The claim was bound to a PersistentVolume and this volume does not exist any longer and all data on it was lost.
conditions	No	Array of PersistentVolumeClaimCondition objects	Current Condition of persistent volume claim. If underlying persistent volume is being resized then the Condition will be set to ResizeStarted.

表 6-300 PersistentVolumeClaimCondition 字段数据结构说明

名称	是否必选	参数类型	描述
type	No	String	Type of the condition. Currently only Ready. Resizing - An user trigger resize of pvc has been started

名称	是否必选	参数类型	描述
status	No	String	Status of the condition. Can be True, False, or Unknown.
lastProbeTime	No	String	Last time we probed the condition.
lastTransitionTime	No	String	Last time the condition transitioned from one status to another.
reason	No	String	Unique, one-word, CamelCase reason for the condition's last transition.
message	No	String	Human-readable message indicating details about last transition.

表 6-301 ResourceRequirements 字段数据结构说明

名称	是否必选	类型	描述
limits	No	Array of ResourceName objects	Maximum amount of compute resources allowed. 说明 limits和requests值必须相等，否则会报错。
requests	No	Array of ResourceName objects	Minimum amount of compute resources required. If Requests is omitted for a container, it defaults to Limits if that is explicitly specified, otherwise to an implementation-defined value. 云容器实例中Pod规格有限制，具体的限制请参见 使用限制 页面的“Pod规格”部分。

表 6-302 ResourceName 字段可选值说明

参数类型	是否可选	参数类型	描述
storage	No	String	Volume size, in bytes (e.g. 5Gi = 5GiB = 5 * 1024 * 1024 * 1024)
cpu	No	String	CPU size, in cores. (500m = .5 cores)
memory	No	String	Memory size, in bytes. (500Gi = 500GiB = 500 * 1024 * 1024 * 1024)

参数类型	是否可选	参数类型	描述
localdir	No	String	Local Storage for LocalDir, in bytes. (500Gi = 500GiB = 500 * 1024 * 1024 * 1024)
nvidia.com/gpu-tesla-v100-16GB	No	String	NVIDIA GPU resource, the type may change in different environments, in production environment is nvidia.com/gpu-tesla-v100-16GB now. The value must be an integer and not less than 1.

表 6-303 labelSelector 字段数据结构说明

名称	是否必选	参数类型	描述
matchExpressions	No	Array of LabelSelector Requirement objects	matchExpressions is a list of label selector requirements. The requirements are ANDed.
matchLabels	No	Object	matchLabels is a map of {key,value} pairs. A single {key,value} in the matchLabels map is equivalent to an element of matchExpressions, whose key field is "key", the operator is "In", and the values array contains only "value". The requirements are ANDed.

表 6-304 LabelSelectorRequirement 字段数据结构说明

名称	是否必选	参数类型	描述
key	No	String	key is the label key that the selector applies to.
operator	No	String	operator represents a key's relationship to a set of values. Valid operators are In, NotIn, Exists and DoesNotExist.

名称	是否必选	参数类型	描述
values	No	Array of strings	values is an array of string values. If the operator is In or NotIn, the values array must be non-empty. If the operator is Exists or DoesNotExist, the values array must be empty. This array is replaced during a strategic merge patch.

表 6-305 PersistentVolume 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	Yes	metadata object	Standard object's metadata.
spec	Yes	spec object	Spec defines a specification of a persistent volume owned by the cluster. Provisioned by an administrator.
status	No	status object	Status represents the current information/status for the persistent volume. Populated by the system. Read-only.

表 6-306 spec 字段数据结构说明

参数	是否必选	参数类型	描述
accessModes	Yes	Array of strings	Access mode. Options: ReadWriteOnce: can be read and written by a single node. ReadOnlyMany: can only be read by multiple nodes. ReadWriteMany: can be read and written by multiple nodes.
capacity	Yes	Object	A description of the persistent volume's resources and capacity.
claimRef	No	claimRef object	ClaimRef is part of a bi-directional binding between PersistentVolume and PersistentVolumeClaim. Expected to be non-nil when bound. claim. VolumeName is the authoritative bind between PV and PVC.
hostPath	No	hostPath object	HostPath represents a directory on the host. Provisioned by a developer or tester. This is useful for single-node development and testing only! On-host storage is not supported in any way and WILL NOT WORK in a multi-node cluster.
nfs	No	nfs object	NFS represents an NFS mount on the host. Provisioned by an admin.
persistentVolumeReclaimPolicy	No	String	What happens to a persistent volume when released from its claim. Valid options are Retain (default) and Recycle. Recycling must be supported by the volume plugin underlying this persistent volume.
storageClassName	No	String	Name of StorageClass to which this persistent volume belongs. Empty value means that this volume does not belong to any StorageClass.

表 6-307 status 字段数据结构说明

参数	是否必选	参数类型	描述
message	No	String	A human-readable message indicating details about why the volume is in this state.
phase	No	String	Phase indicates if a volume is available, bound to a claim, or released by a claim.
reason	No	String	Reason is a brief CamelCase string that describes any failure and is meant for machine parsing and tidy display in the CLI.

表 6-308 claimRef 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	No	String	API version of the referent.
fieldPath	No	String	If referring to a piece of an object instead of an entire object, this string should contain a valid JSON/Go field access statement, such as <code>desiredState.manifest.containers[2]</code> . For example, if the object reference is to a container within a pod, this would take on a value like: <code>"spec.containers{name}"</code> (where "name" refers to the name of the container that triggered the event) or if no container name is specified <code>"spec.containers[2]"</code> (container with index 2 in this pod). This syntax is chosen only to have some well-defined way of referencing a part of an object.
kind	No	String	Kind of the referent.
name	No	String	Name of the referent.
namespace	No	String	Namespace of the referent.
resourceVersion	No	String	Specific resourceVersion to which this reference is made, if any.

参数	是否必选	参数类型	描述
uid	No	String	UID of the referent.

表 6-309 hostPath 字段数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Path of the directory on the host.

表 6-310 nfs 字段数据结构说明

参数	是否必选	参数类型	描述
path	No	String	Path that is exported by the NFS server.
readOnly	No	Boolean	ReadOnly here will force the NFS export to be mounted with read-only permissions. Defaults to false.
server	No	String	Server is the hostname or IP address of the NFS server.

表 6-311 metadata 字段数据结构说明

名称	是否必选	参数类型	描述
name	Yes	String	Name must be unique within a namespace. Is required when creating resources, although some resources may allow a client to request the generation of an appropriate name automatically. Name is primarily intended for creation idempotence and configuration definition. Cannot be updated. 0 characters < name length ≤ 253 characters. The name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.

名称	是否必选	参数类型	描述
clusterName	No	String	The name of the cluster which the object belongs to. This is used to distinguish resources with same name and namespace in different clusters. This field is not set anywhere right now and apiserver is going to ignore it if set in create or update request.
initializers	No	initializers object	An initializer is a controller which enforces some system invariant at object creation time. This field is a list of initializers that have not yet acted on this object. If nil or empty, this object has been completely initialized. Otherwise, the object is considered uninitialized and is hidden (in list/watch and get calls) from clients that haven't explicitly asked to observe uninitialized objects. When an object is created, the system will populate this list with the current set of initializers. Only privileged users may set or modify this list. Once it is empty, it may not be modified further by any user.
enable	No	Boolean	Enable identify whether the resource is available.

名称	是否必选	参数类型	描述
generateName	No	String	<p>An optional prefix used by the server to generate a unique name ONLY IF the Name field has not been provided. If this field is used, the name returned to the client will be different than the name passed. This value will also be combined with a unique suffix. The provided value has the same validation rules as the Name field, and may be truncated by the length of the suffix required to make the value unique on the server.</p> <p>If this field is specified and the generated name exists, the server will NOT return a 409. Instead, it will either return 201 Created or 500 with Reason ServerTimeout indicating a unique name could not be found in the time allotted, and the client should retry (optionally after the time indicated in the Retry-After header).</p> <p>Applied only if Name is not specified.</p> <p>0 characters < generated name length ≤ 253 characters.</p> <p>The generated name must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.</p>
namespace	No	String	<p>Namespace defines the space within each name must be unique. An empty namespace is equivalent to the "default" namespace, but "default" is the canonical representation. Not all objects are required to be scoped to a namespace - the value of this field for those objects will be empty. Must be a DNS_LABEL. Cannot be updated.</p> <p>0 characters < namespace length ≤ 63 characters.</p> <p>The namespace must be a regular expression [a-z0-9]([-a-z0-9]*[a-z0-9])?.</p>

名称	是否必选	参数类型	描述
selfLink	No	String	<p>A URL representing this object. Populated by the system. Read-only.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
uid	No	String	<p>UID is the unique in time and space value for this object. It is typically generated by the server on successful creation of a resource and is not allowed to change on PUT operations. Populated by the system. Read-only.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
resourceVersion	No	String	<p>An opaque value that represents the internal version of this object that can be used by clients to determine when objects have changed. May be used for optimistic concurrency, change detection, and the watch operation on a resource or set of resources. Clients must treat these values as opaque and passed unmodified back to the server. They may only be valid for a particular resource or set of resources. Populated by the system. Read-only. Value must be treated as opaque by clients.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
generation	No	Integer	<p>A sequence number representing a specific generation of the desired state. Currently only implemented by replication controllers. Populated by the system. Read-only.</p>

名称	是否必选	参数类型	描述
creationTimestamp	No	String	<p>A timestamp representing the server time when this object was created. It is not guaranteed to be set in happens-before order across separate operations. Clients may not set this value. It is represented in RFC3339 form and is in UTC. Populated by the system. Read-only. Null for lists.</p> <p>说明 This field is automatically generated. Do not assign any value to this field. Otherwise, API calls would fail.</p>
deletionTimestamp	No	String	<p>RFC 3339 date and time at which this resource will be deleted. This field is set by the server when a graceful deletion is requested by the user, and is not directly settable by a client. The resource will be deleted (no longer visible from resource lists, and not reachable by name) after the time in this field. Once set, this value may not be unset or be set further into the future, although it may be shortened or the resource may be deleted prior to this time. For example, a user may request that a pod is deleted in 30 seconds. The Kubelet will react by sending a graceful termination signal to the containers in the pod. Once the resource is deleted in the API, the Kubelet will send a hard termination signal to the container. If not set, graceful deletion of the object has not been requested. Populated by the system when a graceful deletion is requested. Read-only.</p>
deletionGracePeriodSeconds	No	Integer	<p>Number of seconds allowed for this object to gracefully terminate before it will be removed from the system. Only set when deletionTimestamp is also set. May only be shortened. Read-only.</p>

名称	是否必选	参数类型	描述
labels	Yes	Object	<p>Map of string keys and values that can be used to organize and categorize (scope and select) objects. May match selectors of replication controllers and services.</p> <p>说明 This field should be filled in to create the real storage dynamically. The value of the field is according to the real region and zone.</p>
annotations	No	annotations object	<p>An unstructured key value map stored with a resource that may be set by external tools to store and retrieve arbitrary metadata. They are not queryable and should be preserved when modifying objects.</p> <p>说明 This field should be filled in to create the real storage dynamically. This field indicates the storage plugin and the StorageClass.</p>
ownerReferences	No	ownerReferences object	<p>List of objects depended by this object. If ALL objects in the list have been deleted, this object will be garbage collected. If this object is managed by a controller, then an entry in this list will point to this controller, with the controller field set to true. There cannot be more than one managing controller.</p>
finalizers	No	Array of strings	<p>Must be empty before the object is deleted from the registry. Each entry is an identifier for the responsible component that will remove the entry from the list. If the deletionTimestamp of the object is non-nil, entries in this list can only be removed.</p>

表 6-312 annotations 字段数据结构说明

参数	是否必须	参数类型	描述
volume.beta.kubernetes.io/storage-class	Yes	String	存储类型。 <ul style="list-style-type: none"> 云硬盘存储卷（EVS） 当前支持高I/O（sas）、超高I/O（ssd）和普通I/O（sata） 文件存储卷（SFS） 当前支持标准文件协议类型（nfs-rw）
volume.beta.kubernetes.io/storage-provisioner	Yes	String	挂载路径。 <ul style="list-style-type: none"> 存储类型为EVS时，配置为flexvolume-huawei.com/fuxivol 存储类型为SFS时，配置为flexvolume-huawei.com/fuxinfs

表 6-313 initializers 字段数据结构说明

参数	是否必选	参数类型	描述
pending	No	pending object	Pending is a list of initializers that must execute in order before this object is visible. When the last pending initializer is removed, and no failing result is set, the initializers struct will be set to nil and the object is considered as initialized and visible to all clients.
result	No	result object	If result is set with the Failure field, the object will be persisted to storage and then deleted, ensuring that other clients can observe the deletion.

表 6-314 pending 字段数据结构说明

参数	是否必选	参数类型	描述
name	No	String	name of the process that is responsible for initializing this object.

表 6-315 result 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
code	No	Integer	Suggested HTTP return code for this status, 0 if not set.
details	No	details object	Extended data associated with the reason. Each reason may define its own extended details. This field is optional and the data returned is not guaranteed to conform to any schema except that defined by the reason type.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
message	No	String	A human-readable description of the status of this operation.
metadata	Yes	ListMeta object	Standard list metadata.
reason	No	String	A machine-readable description of why this operation is in the "Failure" status. If this value is empty there is no information available. A Reason clarifies an HTTP status code but does not override it.
status	No	String	Status of the operation. One of: "Success" or "Failure".

表 6-316 details 字段数据结构说明

参数	是否必选	参数类型	描述
causes	No	causes object	The Causes array includes more details associated with the StatusReason failure. Not all StatusReasons may provide detailed causes.
group	No	String	The group attribute of the resource associated with the status StatusReason.
kind	No	String	The kind attribute of the resource associated with the status StatusReason. On some operations may differ from the requested resource Kind.
name	No	String	The name attribute of the resource associated with the status StatusReason (when there is a single name which can be described).
retryAfterSeconds	No	Integer	If specified, the time in seconds before the operation should be retried.
uid	No	String	UID of the resource. (when there is a single resource which can be described).

表 6-317 ListMeta 字段数据结构说明

参数	是否必选	参数类型	描述
resourceVersion	No	String	String that identifies the server's internal version of this object that can be used by clients to determine when objects have changed. Value must be treated as opaque by clients and passed unmodified back to the server. Populated by the system. Read-only.

参数	是否必选	参数类型	描述
Continue	No	String	Continue may be set if the user set a limit on the number of items returned, and indicates that the server has more data available. The value is opaque and may be used to issue another request to the endpoint that served this list to retrieve the next set of available objects. Continuing a list may not be possible if the server configuration has changed or more than a few minutes have passed. The resourceVersion field returned when using this continue value will be identical to the value in the first response.
selfLink	No	String	SelfLink is a URL representing this object. Populated by the system. Read-only.

表 6-318 causes 字段数据结构说明

参数	是否必选	参数类型	描述
field	No	String	The field of the resource that has caused this error, as named by its JSON serialization. May include dot and postfix notation for nested attributes. Arrays are zero-indexed. Fields may appear more than once in an array of causes due to fields having multiple errors. Optional. Examples: "name" - the field "name" on the current resource "items[0].name" - the field "name" on the first array entry in "items"
message	No	String	A human-readable description of the cause of the error. This field may be presented as-is to a reader.

参数	是否必选	参数类型	描述
reason	No	String	A machine-readable description of the cause of the error. If this value is empty there is no information available.

表 6-319 ownerReferences 字段数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	API version of the referent.
blockOwnerDeletion	No	Boolean	If true, AND if the owner has the "foregroundDeletion" finalizer, then the owner cannot be deleted from the key-value store until this reference is removed. Defaults to false. To set this field, a user needs "delete" permission of the owner, otherwise 422 (Unprocessable Entity) will be returned.
kind	Yes	String	Kind of the referent.
name	Yes	String	Name of the referent.
uid	No	String	UID of the referent.
controller	No	Boolean	If true, this reference points to the managing controller.

表 6-320 volume 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.

参数	是否必选	参数类型	描述
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated. In CamelCase.
metadata	Yes	metadata object	Standard object's metadata.
spec	Yes	spec object	Spec defines a specification of a volume owned by the cluster.
status	No	status object	Status represents the current information/status for the volume. Populated by the system. Read-only.

表 6-321 spec 字段数据结构说明

参数	是否必选	参数类型	描述
name	Yes	String	Name of this volume.
size	Yes	Integer	Size of this volume.
description	No	String	Description of this volume.
storageclassname	No	String	StorageClassName new add, use to get az and type from k8s.
inresourcepool	Yes	Boolean	是否在资源池中。
availability_zone	No	String	AvailabilityZone of this volume.
volume_type	No	String	VolumeType of this volume.
snapshot_id	No	String	SnapshotId of this volume.
multiattach	Yes	Boolean	Multiattach defines whether to attach by multiple containers.
storage_type	No	String	Optional values: BS(Block Storage), OS(Object Storage), NFS(Network File System).Default: BS
share_proto	No	String	When storage_type NFS is required, effective value is NFS.

参数	是否必选	参数类型	描述
is_public	No	Boolean	When storage_type is NFS, the visibility of sharing is expressed. Set to true, public visible, set to false, private visible. Default:false
access_to	No	String	When storage_type NFS is required, the definition of the access rule. The length of 1~255, is VPC ID.
access_level	No	String	When storage_type NFS is required, said sharing permission level to access the value of RO (read only), RW (read and write).
pvc_name	No	String	pvcName of volume
access	Yes	Array of access object	sfs access.
vpc_id	Yes	String	efs vpc
enterprise_project_id	No	String	enterprise_project_id
volume_id	No	String	volume_id
auto_expand	No	Boolean	When storage_type NFS is required, when value is true capacity expansion is not supported.

表 6-322 status 字段数据结构说明

参数	是否必选	参数类型	描述
id	No	String	A human-readable message indicating details about why the volume is in this state.
status	Yes	String	Phase indicates if a volume is available, bound to a claim, or released by a claim.
created_at	No	String	Reason is a brief CamelCase string that describes any failure and is meant for machine parsing and tidy display in the CLI.

参数	是否必选	参数类型	描述
attachments	Yes	Array of attachment objects	Attachments information of this volume.
app_info	Yes	Array of app_info objects	volume using info.
access_state	No	String	Access state
access_id	No	String	Vpc id
export_location	No	String	Export location
export_location s	No	String	Export locations

表 6-323 access 字段数据结构说明

参数	是否必选	参数类型	描述
share_id	No	String	uuid of share.
access_type	Yes	String	access rule type.
access_to	Yes	String	vpc id.
access_level	Yes	String	access level.
id	Yes	String	uuid of access rule.
state	Yes	String	status of access rule.should be active or error.

表 6-324 attachment 字段数据结构说明

参数	是否必选	参数类型	描述
attachment_id	No	String	AttachmentId.
server_id	No	String	Server of Attached device.
host_name	No	String	Host name of the attached machine.
device	No	String	Attached device.

表 6-325 app_info 字段数据结构说明

参数	是否必选	参数类型	描述
app_name	Yes	String	App name.
namespace	Yes	String	namespace.
mount_path	Yes	String	Mount path.
app_type	Yes	String	App type

7 权限和授权项

7.1 权限及授权项说明

如果您需要对您所拥有的云容器实例（CCI）进行精细的权限管理，您可以使用统一身份认证服务（Identity and Access Management，简称IAM），如果云账号已经能满足您的要求，不需要创建独立的IAM用户，您可以跳过本章节，不影响您使用CCI服务的其它功能。

默认情况下，新建的IAM用户没有任何权限，您需要将其加入用户组，并给用户组授予策略或角色，才能使用户组中的用户获得相应的权限，这一过程称为授权。授权后，用户就可以基于已有权限对云服务进行操作。关于策略的语法结构及示例，请参见[IAM权限管理说明](#)。

权限根据授权的精细程度，分为[角色](#)和[策略](#)。角色以服务为粒度，是IAM最初提供了一种根据用户的工作职能定义权限的粗粒度授权机制。策略以API接口为粒度进行权限拆分，授权更加精细，可以精确到某个操作、资源和条件，能够满足企业对权限最小化的安全管控要求。

说明

如果您要允许或是禁止某个接口的操作权限，请使用策略。

账号具备所有接口的调用权限，如果使用账号下的IAM用户发起API请求时，该IAM用户必须具备调用该接口所需的权限，否则，API请求将调用失败。每个接口所需要的权限，与各个接口所对应的授权项相对应，只有发起请求的用户被授予授权项所对应的策略，该用户才能成功调用该接口。例如，用户要调用接口来查询Pod，那么这个IAM用户被授予的策略中必须包含允许“CCI:namespaceSubResource:Get”的授权项，该接口才能调用成功。

支持的授权项

策略包含系统策略和自定义策略，如果系统策略不满足授权要求，管理员可以创建自定义策略，并通过给用户组授予自定义策略来进行精细的访问控制。策略支持的操作与API相对应，授权项列表说明如下：

- 权限：允许或拒绝某项操作。
- 对应API接口：自定义策略实际调用的API接口。

- 授权项：自定义策略中支持的Action，在自定义策略中的Action中写入授权项，可以实现授权项对应的权限功能。
- IAM项目(Project)/企业项目(Enterprise Project)：自定义策略的授权范围，包括IAM项目与企业项目。授权范围如果同时支持IAM项目和企业项目，表示此授权项对应的自定义策略，可以在IAM和企业管理两个服务中给用户组授权并生效。如果仅支持IAM项目，不支持企业项目，表示仅能在IAM中给用户组授权并生效，如果在企业管理中授权，则该自定义策略不生效。关于IAM项目与企业项目的区别，详情请参见：[IAM与企业管理的区别](#)。

📖 说明

“√”表示支持，“x”表示暂不支持。

云容器实例（CCI）支持的自定义策略授权项如下所示：

- **Namespace**，Namespace对象管理接口，包括Namespace对象的创建、查询、修改、删除等接口。
- **Pod**，Pod对象管理接口，包括Pod对象的查询接口。
- **Deployment**，Deployment对象管理接口，包括Deployment对象的创建、查询、修改、删除等接口。
- **StatefulSet**，Statefulset对象管理接口，包括Statefulset对象的创建、查询、修改、删除等接口。
- **Job**，Job对象管理接口，包括Job对象的创建、查询、修改、删除等接口。
- **Service**，Service对象管理接口，包括Service对象的创建、查询、修改、删除等接口。
- **Ingress**，Ingress对象管理接口，包括Ingress对象的创建、查询、修改、删除等接口。
- **Network**，Network对象管理接口，包括Network对象的创建、查询、删除等接口。
- **PersistentVolumeClaim**，PersistentVolumeClaim对象管理接口，包括PersistentVolumeClaim对象的创建、查询、修改、删除等接口。
- **ConfigMap**，Configmap对象管理接口，包括Configmap对象的创建、查询、修改、删除等接口。
- **Secret**，Secret对象管理接口，包括Secret对象的创建、查询、修改、删除等接口。
- **ClusterRole**，ClusterRole对象管理接口，包括ClusterRole的查询等接口。
- **RoleBinding**，RoleBinding对象管理接口，包括RoleBinding对象的创建、查询、修改、删除等接口。

7.2 授权项分类

表 7-1 Namespace

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建 Namespace	POST /api/v1/namespaces	CCI:namespace:create	√	√
查询 namespace	GET /api/v1/namespaces/{name}	CCI:namespace:get	√	√
列出 namespaces	GET /api/v1/namespaces	CCI:namespace:list	√	√
删除 namespace	DELETE /api/v1/namespaces/{name}	CCI:namespace:delete	√	√

表 7-2 Pod

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建Pod	POST /api/v1/namespaces/{namespace}/pods	CCI:namespaceSubResource:Create	√	√
查询Pod	GET /api/v1/namespaces/{namespace}/pods/{name}	CCI:namespaceSubResource:Get	√	√
查询指定 Namespace 下所有Pod	GET /api/v1/namespaces/{namespace}/pods	CCI:namespaceSubResource:List	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询Pod状态	GET /api/v1/namespaces/{namespace}/pods/{name}/status	CCI:namespaceSubResource:Get	√	√
查询Pod日志	GET /api/v1/namespaces/{namespace}/pods/{name}/log	CCI:namespaceSubResource:Get	√	√
列出用户所有的Pod	GET /api/v1/pods	cci:namespaceSubResource:List	√	√
替换Pod	PUT /api/v1/namespaces/{namespace}/pods/{name}	CCI:namespaceSubResource:Update	√	√
更新Pod	PATCH /api/v1/namespaces/{namespace}/pods/{name}	CCI:namespaceSubResource:Update	√	√
删除Pod	DELETE /api/v1/namespaces/{namespace}/pods/{name}	CCI:namespaceSubResource:Delete	√	√
删除所有Pod	DELETE /api/v1/namespaces/{namespace}/pods	CCI:namespaceSubResource:Delete	√	√

表 7-3 Deployment

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建 Deployment	POST /apis/apps/v1/namespaces/{namespace}/deployments	<ul style="list-style-type: none"> CCI:namespaceSubResource:Create elb:loadbalancers:create 	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询 Deployment	GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}	CCI:namespaceSubResource:Get	√	√
查询 Namespace下所有 Deployment	GET /apis/apps/v1/namespaces/{namespace}/deployments	CCI:namespaceSubResource:List	√	√
查询 Deployment的状态	GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}/status	CCI:namespaceSubResource:Get	√	√
查询 Deployment伸缩操作	GET /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale	CCI:namespaceSubResource:Get	√	√
列出用户所有的Deployment	GET /apis/apps/v1/deployments	CCI:namespaceSubResource:List	√	√
替换 Deployment	PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}	CCI:namespaceSubResource:Update	√	√
替换 Deployment伸缩操作	PUT /apis/apps/v1/namespaces/{namespace}/deployments/{name}/scale	CCI:namespaceSubResource:Update	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
更新 Deployment	PATCH /apis/ apps/v1/ namespaces/ {namespace}/ deployments/ {name}	CCI:namespaceS ubResource:Upda te	√	√
更新 Deployment伸 缩操作	PATCH /apis/ apps/v1/ namespaces/ {namespace}/ deployments/ {name}/scale	CCI:namespaceS ubResource:Upda te	√	√
删除 Deployment	DELETE /apis/ apps/v1/ namespaces/ {namespace}/ deployments/ {name}	CCI:namespaceS ubResource:Delet e	√	√
删除 Namespace下 所有 Deployment	DELETE /apis/ apps/v1/ namespaces/ {namespace}/ deployments	CCI:namespaceS ubResource:Delet e	√	√

表 7-4 StatefulSet

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建 StatefulSet	POST /apis/ apps/v1/ namespaces/ {namespace}/ statefulsets	<ul style="list-style-type: none"> CCI:namespace SubResource:Cr eate elb:loadbalanc ers:create 	√	√
查询 StatefulSet	GET /apis/apps/v1/ namespaces/ {namespace}/ statefulsets/ {name}	CCI:namespaceSu bResource:Get	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询指定 Namespace 下所有 StatefulSet	GET /apis/apps/v1/namespaces/{namespace}/statefulsets	CCI:namespaceSubResource:List	√	√
查询 StatefulSet 状态	GET /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}/status	CCI:namespaceSubResource:Get	√	√
列出用户所有的 Statefulset	GET /apis/apps/v1/statefulsets	CCI:namespaceSubResource:List	√	√
替换 StatefulSet	PUT /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}	CCI:namespaceSubResource:Update	√	√
更新 StatefulSet	PATCH /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}	CCI:namespaceSubResource:Update	√	√
删除 StatefulSet	DELETE /apis/apps/v1/namespaces/{namespace}/statefulsets/{name}	CCI:namespaceSubResource>Delete	√	√
删除所有 StatefulSet	DELETE /apis/apps/v1/namespaces/{namespace}/statefulsets	CCI:namespaceSubResource>Delete	√	√

表 7-5 Job

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建Job	POST /apis/batch/v1/namespaces/{namespace}/jobs	<ul style="list-style-type: none"> CCI:namespaceSubResource:Create elb:loadbalancers:create 	√	√
查询Job	GET /apis/batch/v1/namespaces/{namespace}/jobs/{name}	CCI:namespaceSubResource:Get	√	√
查询指定Namespace下所有Job	GET /apis/batch/v1/namespaces/{namespace}/jobs	CCI:namespaceSubResource:List	√	√
查询Job状态	GET /apis/batch/v1/namespaces/{namespace}/jobs/{name}/status	CCI:namespaceSubResource:Get	√	√
列出用户所有Job	GET /apis/batch/v1/jobs	CCI:namespaceSubResource:List	√	√
替换Job	PUT /apis/batch/v1/namespaces/{namespace}/jobs/{name}	CCI:namespaceSubResource:Update	√	√
更新Job	PATCH /apis/batch/v1/namespaces/{namespace}/jobs/{name}	CCI:namespaceSubResource:Update	√	√
删除Job	DELETE /apis/batch/v1/namespaces/{namespace}/jobs/{name}	CCI:namespaceSubResource:Delete	√	√
删除所有Job	DELETE /apis/batch/v1/namespaces/{namespace}/jobs	CCI:namespaceSubResource:Delete	√	√

表 7-6 Service

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建Service	POST /api/v1/namespaces/{namespace}/services	<ul style="list-style-type: none"> CCI:namespaceSubResource:Create elb:loadbalancers:create 	√	√
查询Service	GET /api/v1/namespaces/{namespace}/services/{name}	CCI:namespaceSubResource:Get	√	√
查询所有Service	GET /api/v1/namespaces/{namespace}/services	CCI:namespaceSubResource:List	√	√
查询Service状态	GET /api/v1/namespaces/{namespace}/services/{name}/status	CCI:namespaceSubResource:Get	√	√
删除Service	DELETE /api/v1/namespaces/{namespace}/services/{name}	CCI:namespaceSubResource:Delete	√	√

表 7-7 Ingress

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建Ingress	POST /apis/extensions/v1beta1/namespaces/{namespace}/ingresses	<ul style="list-style-type: none"> CCI:namespaceSubResource:Create elb:loadbalancers:create 	√	√
查询Ingress	GET /apis/extensions/v1beta1/namespaces/{namespace}/ingresses/{name}	CCI:namespaceSubResource:Get	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询所有 Ingress	GET /apis/ extensions/v1beta1/ namespaces/ {namespace}/ ingresses	CCI:namespaceSubResource:List	√	√
查询 Ingress状态	GET /apis/ extensions/v1beta1/ namespaces/ {namespace}/ ingresses/{name}/ status	CCI:namespaceSubResource:Get	√	√
替换 Ingress	PUT /apis/ extensions/v1beta1/ namespaces/ {namespace}/ ingresses/{name}	CCI:namespaceSubResource:Update	√	√
更新 Ingress	PATCH /apis/ extensions/v1beta1/ namespaces/ {namespace}/ ingresses/{name}	CCI:namespaceSubResource:Update	√	√
删除 Ingress	DELETE /apis/ extensions/v1beta1/ namespaces/ {namespace}/ ingresses/{name}	CCI:namespaceSubResource>Delete	√	√
删除所有 Ingress	DELETE /apis/ extensions/v1beta1/ namespaces/ {namespace}/ ingresses	CCI:namespaceSubResource>Delete	√	√

表 7-8 Network

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建Network	POST /apis/networking.cci.io/v1beta1/namespaces/{namespace}/networks	<ul style="list-style-type: none"> • CCI:networkk:Create • vpc:vpcs:create • vpc:ports:create • vpc:vpcs:get • vpc:subnets:get • vpc:publicips:get • vpc:bandwidths:get • vpc:ports:get • vpc:peerings:get • vpc:quotas:list • vpc:privateips:get • vpc:securityGroups:get • vpc:securityGroupRules:get • vpc:networks:get • vpc:routers:get • vpc:floatingips:get • vpc:firewallRules:get 	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询Network	GET /apis/ networking.cci.io /v1beta1/ namespaces/ {namespace}/ networks/ {name}	CCI:network: Get	√	√
查询所有 Network	GET /apis/ networking.cci.io /v1beta1/ namespaces/ {namespace}/ networks	CCI:network:L ist	√	√
查询Network状 态	GET /apis/ networking.cci.io /v1beta1/ namespaces/ {namespace}/ networks/ {name}/status	CCI:network: Get	√	√
删除Network	DELETE /apis/ networking.cci.io /v1beta1/ namespaces/ {namespace}/ networks/ {name}	<ul style="list-style-type: none">• CCI:networ k:Delete• vpc:vpcs:de lete• vpc:ports:d elete	√	√
删除所有 Network对象	DELETE /apis/ networking.cci.io /v1beta1/ namespaces/ {namespace}/ networks	<ul style="list-style-type: none">• CCI:networ k:Delete• vpc:vpcs:de lete• vpc:ports:d elete	√	√

表 7-9 PersistentVolumeClaim

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建 PersistentVolumeClaim	POST /api/v1/namespaces/{namespace}/persistentvolumeclaims	<ul style="list-style-type: none"> • CCI:namespaceSubResource:Create - 云硬盘存储卷 evs:volumes:create evs:volumes:get evs:types:get - 文件存储卷 sfs:shares:createShare sfs:shares:getOSQuotaSets sfs:shares:ShareAction 	√	√
查询 PersistentVolumeClaim	GET /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}	<ul style="list-style-type: none"> • CCI:namespaceSubResource:Get - 云硬盘存储卷 evs:volumes:get - 文件存储卷 sfs:shares:getAllSharesDetail 	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询所有 PersistentVolumeClaim	GET /api/v1/namespaces/{namespace}/persistentvolumeclaims	<ul style="list-style-type: none"> CCI:namespaceSubResource:List 云硬盘存储卷 evs:volume: list 文件存储卷 sfs:shares:getAllSharesDetail sfs:shares:ShareAction 对象存储卷 obs:bucket: ListAllMyBuckets 	√	√
删除 PersistentVolumeClaim	DELETE /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}	<ul style="list-style-type: none"> CCI:namespaceSubResource:Delete 云硬盘存储卷 evs:volume: delete evs:volume: get 文件存储卷 sfs:shares:deleteShare 	√	√

表 7-10 ConfigMap

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建ConfigMap	POST /api/v1/namespaces/{namespace}/configmaps	CCI:namespaceSubResource:Create	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
查询ConfigMap	GET /api/v1/namespaces/{namespace}/configmaps/{name}	CCI:namespaceSubResource:Get	√	√
查询所有ConfigMap	GET /api/v1/namespaces/{namespace}/configmaps	CCI:namespaceSubResource:List	√	√
替换ConfigMap	PUT /api/v1/namespaces/{namespace}/configmaps/{name}	CCI:namespaceSubResource:Update	√	√
更新ConfigMap	PATCH /api/v1/namespaces/{namespace}/configmaps/{name}	CCI:namespaceSubResource:Update	√	√
删除ConfigMap	DELETE /api/v1/namespaces/{namespace}/configmaps/{name}	CCI:namespaceSubResource:Delete	√	√
删除所有ConfigMap	DELETE /api/v1/namespaces/{namespace}/configmaps	CCI:namespaceSubResource:Delete	√	√

表 7-11 Secret

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建Secret	POST /api/v1/namespaces/{namespace}/secrets	CCI:namespaceSubResource:Create	√	√
替换Secret	PUT /api/v1/namespaces/{namespace}/secrets/{name}	CCI:namespaceSubResource:Update	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
更新Secret	PATCH /api/v1/namespaces/{namespace}/secrets/{name}	CCl:namespaceSubResource:Update	√	√
删除Secret	DELETE /api/v1/namespaces/{namespace}/secrets/{name}	CCl:namespaceSubResource:Delete	√	√
删除所有Secret	DELETE /api/v1/namespaces/{namespace}/secrets	CCl:namespaceSubResource:Delete	√	√

表 7-12 ClusterRole

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
获取指定的ClusterRole	GET /apis/rbac.authorization.k8s.io/v1/clusterroles/{name}	CCl:rbac:Get	√	√
获取ClusterRole列表	GET /apis/rbac.authorization.k8s.io/v1/clusterroles	CCl:rbac:List	√	√

表 7-13 RoleBinding

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
创建RoleBinding	POST /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings	CCl:rbac:Create	√	√

权限	对应API接口	授权项	IAM项目 (Project)	企业项目 (Enterprise Project)
更新指定的 RoleBinding	PATCH /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}	CCI:rbac:Update	√	√
替换指定的 RoleBinding	PUT /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}	CCI:rbac:Update	√	√
删除指定的 RoleBinding	DELETE /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}	CCI:rbac:Delete	√	√
获取指定的 RoleBinding	GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings/{name}	CCI:rbac:Get	√	√
获取指定 namespace 下 RoleBinding 列表	GET /apis/rbac.authorization.k8s.io/v1/namespaces/{namespace}/rolebindings	CCI:rbac:List	√	√
获取 RoleBinding 列表	GET /apis/rbac.authorization.k8s.io/v1/rolebindings	CCI:rbac:List	√	√

8 附录

8.1 PATCH 请求方法操作说明

对于PATCH请求方法的操作，Kubernetes API通过相应的HTTP头域“Content-Type”对其进行识别。

操作说明

目前支持两种类型的PATCH请求方法的操作。

1. Merge Patch, Content-Type: application/merge-patch+json

在RFC7386协议的定义中，Merge Patch必须包含对一个资源对象的部分描述，即为JSON对象。该JSON对象被提交到服务端后与服务端的当前对象合并，即替换当前资源对象中的列表域，从而创建一个新的对象。

2. Strategic Merge Patch, Content-Type: application/strategic-merge-patch+json

Strategic Merge Patch是添加合法的元数据到API对象中，并通过这些新的元数据来决定哪个列表被合并，哪个列表不该被合并。当前这些元数据则作为结构标签。

“Merge Patch”与“Strategic Merge Patch”的操作区别请参见[PATCH请求方法操作示例](#)。

8.2 PATCH 请求方法操作示例

本章节主要介绍PATCH请求方法中，“Merge Patch”与“Strategic Merge Patch”类型的操作示例。

操作示例

以创建Deployment资源对象为例。

请求示例

```
{  
  "apiVersion": "v1",  
  "kind": "Deployment",
```

```
"metadata": {
  "name": "nginx"
},
"spec": {
  "replicas": 2,
  "selector": {
    "app": "nginx"
  },
  "template": {
    "metadata": {
      "labels": {
        "app": "nginx"
      }
    },
    "spec": {
      "containers": [
        {
          "name": "redis",
          "image": "redis:latest",
          "ports": [
            {
              "containerPort": 80
            }
          ]
        }
      ]
    }
  }
}
}
```

响应示例

```
{
  "kind": "ReplicationController",
  "apiVersion": "v1",
  "metadata": {
    "name": "frontend-controller",
    "namespace": "default",
    "selfLink": "/api/v1/namespaces/default/replicationcontrollers/nginx-controller",
    "uid": "549b2234-5d46-11e6-aeb9-286ed488fafa",
    "resourceVersion": "4110",
    "generation": 1,
    "creationTimestamp": "2016-08-08T08:58:52Z",
    "labels": {
      "app": "nginx"
    }
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "app": "nginx"
    },
    "template": {
      "metadata": {
        "creationTimestamp": null,
        "labels": {
          "app": "nginx"
        }
      },
      "spec": {
        "containers": [
          {
            "name": "redis",
            "image": "redis:latest",
            "ports": [
              {
                "containerPort": 80,
                "protocol": "TCP"
              }
            ]
          }
        ]
      }
    }
  }
}
```

```
    ],
    "resources": {},
    "terminationMessagePath": "/dev/termination-log",
    "imagePullPolicy": "Always"
  }
],
"restartPolicy": "Always",
"dnsPolicy": "ClusterFirst",
"securityContext": {}
}
}
},
"status": {
  "replicas": 0
}
}
```

- 如果使用“Merge Patch”类型操作，添加一个容器到指定Deployment的“template”参数中，则其中的整个容器列表将被新添加的容器所替换。

Merge Patch请求

```
{
  "spec": {
    "template": {
      "spec": {
        "containers": [
          {
            "name": "hello-world",
            "image": "busybox:latest"
          }
        ]
      }
    }
  }
}
```

Merge Patch响应

```
{
  "kind": "ReplicationController",
  "apiVersion": "v1",
  "metadata": {
    "name": "frontend-controller",
    "namespace": "default",
    "selfLink": "/api/v1/namespaces/default/replicationcontrollers/nginx-controller",
    "uid": "549b2234-5d46-11e6-aeb9-286ed488fafa",
    "resourceVersion": "4159",
    "generation": 2,
    "creationTimestamp": "2016-08-08T08:58:52Z",
    "labels": {
      "app": "nginx"
    }
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "app": "nginx"
    },
    "template": {
      "metadata": {
        "creationTimestamp": null,
        "labels": {
          "app": "nginx"
        }
      },
      "spec": {
        "containers": [
          {
            "name": "hello-world",
            "image": "busybox:latest",
            "resources": {},

```

```
        "terminationMessagePath": "/dev/termination-log",
        "imagePullPolicy": "Always"
      }
    ],
    "restartPolicy": "Always",
    "dnsPolicy": "ClusterFirst",
    "securityContext": {}
  }
},
"status": {
  "replicas": 2,
  "fullyLabeledReplicas": 2,
  "observedGeneration": 1
}
}
```

其中“containers”参数列表被新添加的内容所替换。

- 而使用“Strategic Merge Patch”类型操作，是添加元数据到资源对象中，并通过这些新元数据来决定各个列表是否需要被合并。

Strategic Merge Patch请求

```
{
  "spec": {
    "template": {
      "spec": {
        "containers": [
          {
            "name": "hello-world",
            "image": "busybox:latest"
          }
        ]
      }
    }
  }
}
```

Strategic Merge Patch响应

```
{
  "kind": "ReplicationController",
  "apiVersion": "v1",
  "metadata": {
    "name": "frontend-controller",
    "namespace": "default",
    "selfLink": "/api/v1/namespaces/default/replicationcontrollers/nginx-controller",
    "uid": "f2e048bb-5d46-11e6-aeb9-286ed488fafa",
    "resourceVersion": "4250",
    "generation": 2,
    "creationTimestamp": "2016-08-08T09:03:18Z",
    "labels": {
      "app": "nginx"
    }
  },
  "spec": {
    "replicas": 2,
    "selector": {
      "app": "nginx"
    },
    "template": {
      "metadata": {
        "creationTimestamp": null,
        "labels": {
          "app": "nginx"
        }
      },
      "spec": {
        "containers": [
          {
            "name": "redis",
```

```
        "image": "redis:latest",
        "ports": [
          {
            "containerPort": 80,
            "protocol": "TCP"
          }
        ],
        "resources": {},
        "terminationMessagePath": "/dev/termination-log",
        "imagePullPolicy": "Always"
      },
      {
        "name": "hello-world",
        "image": "busybox:latest",
        "resources": {},
        "terminationMessagePath": "/dev/termination-log",
        "imagePullPolicy": "Always"
      }
    ],
    "restartPolicy": "Always",
    "dnsPolicy": "ClusterFirst",
    "securityContext": {}
  }
},
"status": {
  "replicas": 2,
  "fullyLabeledReplicas": 2,
  "observedGeneration": 1
}
```

其中“containers”参数列表与新添加的内容合并，而不是替换，合并的依据为“name”域的值。

8.3 约束限制

区域

云容器实例当前支持的区域，如表8-1所示。

表 8-1 支持的区域

区域名称	区域
拉美-圣保罗一	sa-brazil-1

Pod 规格

当不使用GPU时，Pod规格需满足如下要求：

- Pod的CPU取值范围为0.25核-32核，或者自定义选择48核、64核，且单个容器的CPU必须为0.25核的整数倍
- Pod的内存取值范围为1GB-512GB，且内存必须为1GB的整数倍
- Pod的CPU/内存配比值必须在1:2到1:8之间
- 一个Pod内最多支持5个容器，单个容器最小配置是0.25核、0.2GB，最大同容器实例的最大配置

- Pod中所有容器和InitContainer（启动容器）规格中的request和limit相等
详情请参见[Pod规格计算方式](#)。

Pod 存储空间

如果没有挂载EVS等磁盘，应用数据存储在不容器的物理机磁盘，每个Pod存储空间限制为CPU物理机磁盘为20G，GPU物理机磁盘为20G。

配额

云容器实例对单个用户的资源数量和容量限定了配额，配额的详细信息请参见[关于配额](#)。

Pod 规格计算方式

Pod规格的计算方式遵循如下规则：

- 步骤1** Pod包含的所有InitContainer上定义的任何特定资源的约束值 (limit) 或请求值 (request) 的最大值，作为Pod有效初始request/limit。
- 步骤2** Pod对资源的有效limit/request，是取如下两项的较大者：
 - 所有应用容器对某个资源的limit/request之和；
 - 对某个资源的有效初始的limit/request。

----结束

说明

InitContainer是一种特殊容器，在 Pod 内的应用容器启动之前运行。有关InitContainer更多解释请参见[对容器进行初始化操作](#)。

8.4 Namespace 和 Network

Namespace（命名空间）是一种在多个用户之间划分资源的方法。适用于用户中存在多个团队或项目的情况。当前云容器实例提供“通用计算型”类型的资源，创建命名空间时需要选择资源类型，后续创建的负载中容器就运行在此类型的集群上。

- 通用计算型**：支持创建含CPU资源的容器实例及工作负载，适用于通用计算场景。

Network是云容器实例扩展的一种Kubernetes资源对象，用于关联VPC及子网，从而使得容器实例能够使用公有云的网络资源。

Namespace 与网络的关系

从网络角度看，命名空间对应一个虚拟私有云（VPC）中一个子网，如[图8-1](#)所示，在创建命名空间时会关联已有VPC或创建一个新的VPC，并在VPC下创建一个子网。后续在该命名空间下创建Pod、Service等资源时都会在对应的VPC及子网之内，且占用子网中的IP地址。

通常情况下，如果您在同一个VPC下还会使用其他服务的资源，您需要考虑您的网络规划，如子网网段划分、IP数量规划等，确保有可用的网络资源。

图 8-1 命名空间与 VPC 子网的关系



哪些情况下适合使用多个命名空间

因为Namespace可以实现部分的环境隔离，当您的项目和人员众多的时候可以考虑根据项目属性，例如生产、测试、开发划分不同的Namespace。

创建 Namespace

Namespace下需要有一个Network关联VPC及子网，创建完Namespace后需要创建一个Network。

说明

通常情况下，没有频繁创建Namespace的需求，建议通过云容器实例的控制台界面创建Namespace，具体方法请参见[创建命名空间](#)。

以下示例创建一个名为namespace-test的Namespace，指定云容器实例的资源类型为general-computing。

```
apiVersion: v1
kind: Namespace
metadata:
  name: namespace-test
  labels:
    sys_enterprise_project_id: "0"
  annotations:
    namespace.kubernetes.io/flavor: general-computing
spec:
  finalizers:
    - kubernetes
```

这里的定义文件采用YAML格式描述（如果您对YAML格式不了解，可以参考[YAML语法](#)），也是使用JSON格式。

- `sys_enterprise_project_id`：表示企业项目ID，可进入[企业管理](#)的企业项目详情页面获取。未开通企业管理的用户无需配置此参数。不配置时默认为0，表示default企业项目。
- `namespace.kubernetes.io/flavor: general-computing`：指定命名空间类型。命名空间的类型有如下两种：
 - **general-computing**：通用计算型，支持创建含CPU资源的容器实例及工作负载，适用于通用计算场景。
 - **gpu-accelerated**：GPU型，支持创建含GPU资源的容器实例及工作负载，适用于深度学习、科学计算、视频处理等场景。

假如上面Namespace定义的文件名称为ns.yaml，则执行kubect create -f ns.yaml即可创建命名空间，-f 表示从文件创建。

```
# kubectl create -f ns.yaml
namespace/namespace-test created
```

执行kubectl get ns查询namespace是否创建成功，ns为namespace的缩写。

```
# kubectl get ns
NAME          STATUS   AGE
namespace-test Active   23s
```

如上，可以看到namespace-test这个命名空间创建成功，且存在的时长为23秒。

登录云容器实例控制台，单击左侧导航栏“命名空间”，您可以看到命名空间创建成功，但状态为“异常”。这是因为在云容器实例中，您需要为Namespace定义网络策略，具体操作方法请参见[创建Network](#)。

图 8-2 Namespace-异常

名称	类型	状态
namespace-test	通用计算型	 异常

创建 Network

Namespace创建好后，需要为Namespace创建网络策略，关联VPC及子网。

以下示例创建一个名为test-network的Network。

```
apiVersion: networking.cci.io/v1beta1
kind: Network
metadata:
  annotations:
    network.alpha.kubernetes.io/default-security-group: security-group-id
    network.alpha.kubernetes.io/domain-id: domain-id
    network.alpha.kubernetes.io/project-id: project-id
  name: test-network
spec:
  cidr: 192.168.0.0/24
  attachedVPC: vpc-id
  networkID: network-id
  networkType: underlay_neutron
  subnetID: subnet-id
```

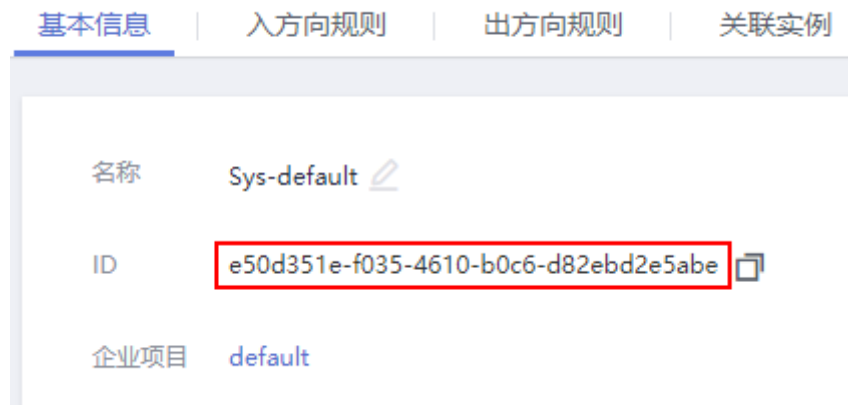
📖 说明

此处VPC和子网的网段不能为10.247.0.0/16，10.247.0.0/16是云容器实例预留给Service的网段。如果您使用此网段，后续可能会造成IP冲突，导致负载无法创建或服务不可用；如果您不需要通过Service访问，而是直接访问Pod，则可以使用此网段。

上面参数获取方法如下：

- network.alpha.kubernetes.io/domain-id：账号ID，可以在[我的凭证](#)获取。
- network.alpha.kubernetes.io/project-id：项目ID，可以在[我的凭证](#)获取。
- network.alpha.kubernetes.io/default-security-group：安全组ID，可以在[安全组控制台](#)获取，如下图。

图 8-3 获取安全组 ID



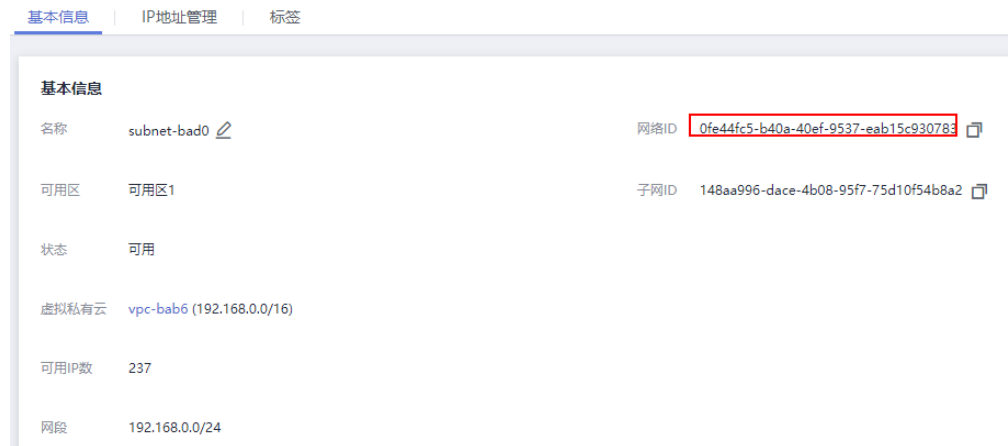
- cidr: 子网网段。
- attachedVPC: 虚拟私有云的ID, 可在VPC控制台获取。

图 8-4 获取 VPC ID



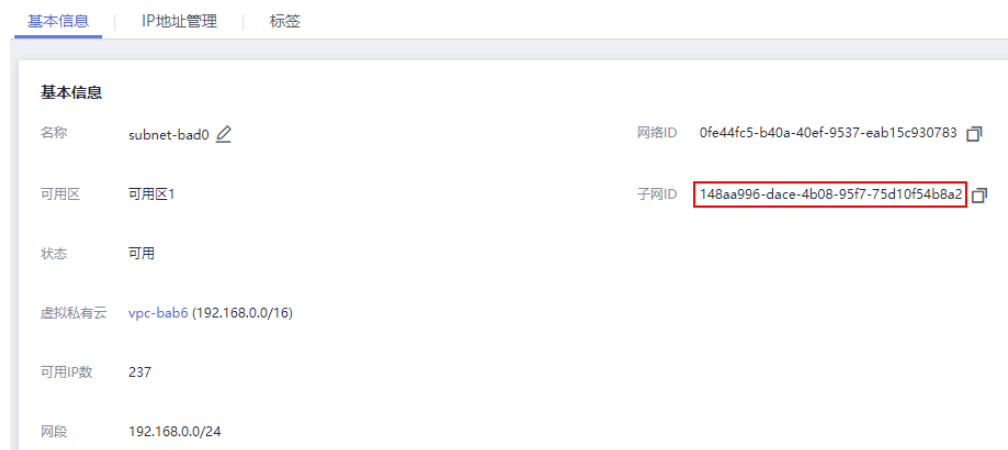
- networkID: 子网的网络ID, 可在VPC控制台 > 子网中获取。

图 8-5 获取子网的网络 ID



- networkType: 网络类型，当前仅支持underlay_neutron网络模式。
- subnetID: 子网ID，可在VPC控制台 > 子网获取。

图 8-6 获取子网 ID



假如上面Network定义的文件名称为network.yaml，则执行kubect create -f network.yaml即可创建命名空间，-f 表示从文件创建。这里--namespace namespace-test是指定在namespace-test这个命名空间下创建。

```
# kubectl create -f network.yaml --namespace namespace-test
network.networking.cci.io/test-network created
```

登录云容器实例控制台，单击左侧导航栏“命名空间”，您可以看到命名空间创建成功，且状态为“正常”。

图 8-7 Namespace-正常

名称	类型	状态
namespace-test	通用计算型	正常

为 kubectl 上下文指定 Namespace

上面创建Network是在指定的Namespace下创建的，本文档后续的资源创建都是在某个命名空间下操作，每次都指定命名空间比较麻烦，您可以为kubectl上下文指定命名空间，这样在某个上下文中，创建的资源就都是在某个命名空间下，方便操作。

指定Namespace只需要在设置上下文命令中添加一个“--namespace”选项，如下所示。

```
kubectl config set-context $context --namespace=$ns
```

其中，*\$ns*为Namespace的名称；*\$context*为上下文的名称，可以自定义，也可执行如下命令获取：

```
# kubectl config get-contexts
CURRENT_NAME          CLUSTER          AUTHINFO
NAMESPACE
  cci-context-cn-east-3-1C8PNI0POPPCSFGXPM6S  cci-cluster-cn-east-3  cci-user-cn-east-3-1C8PNI0POPPCSFGXPM6S
* cci-context-cn-east-3-hwuser_xxx            cci-cluster-cn-east-3  cci-user-cn-east-3-hwuser_xxx
  kubernetes-admin@kubernetes                kubernetes            kubernetes-admin
```

假设，上面创建的Namespace名称为namespace-test，则示例如下。

```
# kubectl config set-context cci-context --namespace=namespace-test
```

指定Namespace后，就可以使用 kubectl 命令直接操作云容器实例的相关资源。如下所示，执行**kubectl get pod**，查看Pod资源，一切正常。

```
# kubectl get pod
No resources found.
```

8.5 状态码

状态码	编码	状态说明
100	Continue	继续请求。 这个临时响应用来通知客户端，它的部分请求已经被服务器接收，且仍未被拒绝。
101	Switching Protocols	切换协议。只能切换到更高级的协议。 例如，切换到HTTP的新版本协议。
200	OK	GET、PUT、POST操作正常返回。
201	Created	创建类的请求完全成功。
202	Accepted	已经接受请求，但未处理完成。
203	Non-Authoritative Information	非授权信息，请求成功。
204	NoContent	请求完全成功，同时HTTP响应不包含响应体。 在响应OPTIONS方法的HTTP请求时返回此状态码。
205	Reset Content	重置内容，服务器处理成功。

状态码	编码	状态说明
206	Partial Content	服务器成功处理了部分GET请求。
300	Multiple Choices	多种选择。请求的资源可包括多个位置，相应可返回一个资源特征与地址的列表用于用户终端（例如：浏览器）选择。
301	Moved Permanently	永久移动，请求的资源已被永久的移动到新的URI，返回信息会包括新的URI。
302	Found	资源被临时移动。
303	See Other	查看其它地址。 使用GET和POST请求查看。
304	Not Modified	所请求的资源未修改，服务器返回此状态码时，不会返回任何资源。
305	Use Proxy	所请求的资源必须通过代理访问。
306	Unused	已经被废弃的HTTP状态码。
400	BadRequest	非法请求。 建议直接修改该请求，不要重试该请求。
401	Unauthorized	在客户端提供认证信息后，返回该状态码，表明服务端指出客户端所提供的认证信息不正确或非法。
402	Payment Required	保留请求。
403	Forbidden	请求被拒绝访问。 返回该状态码，表明请求能够到达服务端，且服务端能够理解用户请求，但是拒绝做更多的事情，因为该请求被设置为拒绝访问，建议直接修改该请求，不要重试该请求。
404	NotFound	所请求的资源不存在。 建议直接修改该请求，不要重试该请求。
405	MethodNotAllowed	请求中带有该资源不支持的方法。 建议直接修改该请求，不要重试该请求。
406	Not Acceptable	服务器无法根据客户端请求的内容特性完成请求。
407	Proxy Authentication Required	请求要求代理的身份认证，与401类似，但请求者应当使用代理进行授权。
408	Request Time-out	服务器等候请求时发生超时。 客户端可以随时再次提交该请求而无需进行任何更改。

状态码	编码	状态说明
409	Conflict	服务器在完成请求时发生冲突。 返回该状态码，表明客户端尝试创建的资源已经存在，或者由于冲突请求的更新操作不能被完成。
410	Gone	客户端请求的资源已经不存在。 返回该状态码，表明请求的资源已被永久删除。
411	Length Required	服务器无法处理客户端发送的不带Content-Length的请求信息。
412	Precondition Failed	未满足前提条件，服务器未满足请求者在请求中设置的其中一个前提条件。
413	Request Entity Too Large	由于请求的实体过大，服务器无法处理，因此拒绝请求。为防止客户端的连续请求，服务器可能会关闭连接。如果只是服务器暂时无法处理，则会包含一个Retry-After的响应信息。
414	Request-URI Too Large	请求的URI过长（URI通常为网址），服务器无法处理。
415	Unsupported Media Type	服务器无法处理请求附带的媒体格式。
416	Requested range not satisfiable	客户端请求的范围无效。
417	Expectation Failed	服务器无法满足Expect的请求头信息。
422	Unprocessable Entity	请求格式正确，但是由于含有语义错误，无法响应。
429	TooManyRequests	表明请求超出了客户端访问频率的限制或者服务端接收到多于它能处理的请求。建议客户端读取相应的Retry-After首部，然后等待该首部指出的时间后再重试。
500	InternalServerError	表明服务端能被请求访问到，但是不能理解用户的请求。
501	Not Implemented	服务器不支持请求的功能，无法完成请求。
502	Bad Gateway	充当网关或代理的服务器，从远端服务器接收到了一个无效的请求。
503	ServiceUnavailable	被请求的服务无效。 建议直接修改该请求，不要重试该请求。
504	ServerTimeout	请求在给定的时间内无法完成。客户端仅在为请求指定超时（Timeout）参数时会得到该响应。
505	HTTP Version not supported	服务器不支持请求的HTTP协议的版本，无法完成处理。

8.6 错误码

调用接口出错后，将不会返回结果数据。调用方可根据每个接口对应的错误码来定位错误原因。当调用出错时，HTTP请求返回一个 4xx 或 5xx 的HTTP状态码。返回的消息体中是具体的错误代码及错误信息。在调用方找不到错误原因时，可以联系华为云客服，并提供错误码，以便我们尽快帮您解决问题。

错误响应 Body 体格式说明

当接口调用出错时，会返回错误码及错误信息说明，错误响应的Body体格式如下所示。

```
{
  "kind": "Status",
  "apiVersion": "v1",
  "metadata": {},
  "status": "Failure",
  "message": "namespace name is already exist",
  "reason": "Conflict",
  "code": 409,
  "errorCode": "CCI.02.409101"
}
```

其中，code表示HTTP状态码，errorCode表示错误码，message表示错误描述信息。

错误码说明

当您使用云容器实例的API时，如果遇到“APIGW”开头的错误码，请参见[API网关错误码](#)进行处理。

状态码	错误码	错误信息	描述	处理措施
400	CCI.01.400101	request body error	请求错误: 请求体不合法	按接口文档排查请求体
400	CCI.01.400102	provide namespace request ed	查询 namespace 详情失败: 未指定待查询的 namespace	请指定待查询的namespace，请修改后重试
400	CCI.01.400103	missing user token in the request header	创建网络失败: 请求头未携带最终用户的 token:X-User-Token	请求头携带最终用户的token:X-User-Token

状态码	错误码	错误信息	描述	处理措施
400	CCI.01.400104	no agency quota found from iam, please enlarge your quota	委托配额不足, 请到配额中心申请调整配额	请到配额中心申请调整配额
400	CCI.01.400105	must specify namespace name through fieldSelector, example: fieldSelector=metadata.name=namespaceName	必须通过fieldSelector来指定namespace名称, 例如: fieldSelector=metadata.name=namespaceName	请通过fieldSelector来指定namespace名称
400	CCI.01.400106	volume import:request body error	卷导入接口的请求体不合法	按卷导入接口文档排查请求体
400	CCI.01.400107	enterprise id not valid	企业项目id不合法	企业项目id不合法, 请修改后重试

状态码	错误码	错误信息	描述	处理措施
400	CCI.01.400201	subnet cidr subnet invalid : within the coverage of VPC, and no more subnets overlap	子网cidr非法: 在vpc的覆盖范围内, 和其他的子网重叠	修改子网网段, 不与其他子网重叠
400	CCI.01.400202	project id in network request body invalid	创建网络失败: 请求体的 Annotation 中未携带 project_id	请求体的Annotation中需携带project_id, 请修改后重试
400	CCI.01.400301	storage type parameter invalid : nfs/obs/bs	存储类型参数错误, 当前支持的存储类型有:nfs、obs和bs	存储类型参数错误, 请修改后重试

状态码	错误码	错误信息	描述	处理措施
400	CCI.01.400302	storage volume id format invalid: the beginning and end of a numeric letter, which may contain an underline in the middle of the numeric letter	存储卷id格式非法: 数字字母的开头和结尾, 数字字母的中间可以包含下划线	请修改存储卷id格式后重试
400	CCI.01.400303	storage volume id len invalid : no bigger than 64	存储卷id长度非法: 最大长度为64	存储卷id长度最大为64, 请修改后重试
400	CCI.01.400304	storage request body invalid	释放存储卷失败: delete_volume字段不合法	请修改delete_volume字段后重试
400	CCI.01.400305	storage volume not available	导入存储卷失败: 存储卷状态异常	联系技术支持排查存储卷问题后重试
400	CCI.01.400306	storage volume is not found	导入存储卷失败: 指定的存储卷不存在	请确定导入的存储卷存在后尝试

状态码	错误码	错误信息	描述	处理措施
400	CCI.01.400307	nfs storage class type err,use "nfs-rw"	导入存储卷失败：指定的nfs存储的storageclass不支持，当前支持：nfs-rw	请导入nfs-rw
400	CCI.01.400308	storage volume type err ,available: sas/ssd/sata	导入存储卷失败：存储规格不支持，支持的规格为:sas、ssd和sata	请导入支持的存储格式，如sas、ssd和sata
400	CCI.01.400309	storage volume already provisioned	导入存储卷失败：当前卷已经被导入	请确认当前卷是否已经被导入
400	CCI.01.400401	invalid billing request body: unsupported billing status	billing接口失败：不支持的status字段	请修改支持的status字段
400	CCI.01.400402	invalid billing request body: unsupported billing resource type	billing接口失败：不支持的resource type字段	请修改支持的resource type字段

状态码	错误码	错误信息	描述	处理措施
400	CCI.01.400403	invalid billing request user: inner user can not be freeze or unfreeze	billing接口失败: 不支持操作内部租户	修改user为非内部用户
400	CCI.01.400404	invalid billing request body: resource infos can not match scene	billing接口失败: 不支持的scene字段	请修改支持的scene字段
500	CCI.01.500101	find no cluster endpoint for namespace xxx from icluster response header	通过 namespace 查询集群地址失败	联系技术支持解决
403	CCI.01.403009	your account is in arrears, please recharge	请求错误: 账户欠费	请先给账户充值
403	CCI.01.403101	only gpu beta users can use gpu in cci	请求出错: 未申请GPU公测	请先申请GPU公测

状态码	错误码	错误信息	描述	处理措施
403	CCI.01.403102	only obs beta users can use obs in cci	请求出错：未申请OBS公测	请先申请OBS公测
403	CCI.01.403103	user's token can not match auth token for resourceUser's	鉴权出错：token中domain与请求domain不匹配	请确认token中domain与请求domain相匹配
403	CCI.01.403104	current user has no right	鉴权出错：用户无权限进行当前操作	用户无权限进行当前操作，请确认当前用户是否有操作权限
403	CCI.01.403105	insufficient operation permission, require cci:namespace:create	创建命名空间失败，您可能没有cci:namespace:create权限	请向管理员申请namespace create权限
403	CCI.01.403106	insufficient operation permission, require cci:namespace:delete	删除命名空间失败，您可能没有cci:namespace:delete权限	请向管理员申请namespace delete权限

状态码	错误码	错误信息	描述	处理措施
403	CCI.01.403107	insufficient operation permission, require cci:namespace:update	更新命名空间失败, 您可能没有 cci:namespace:update 权限	请向管理员申请 namespace update 权限
403	CCI.01.403108	insufficient operation permission, require cci:namespace:get	获取命名空间失败, 您可能没有 cci:namespace:get 权限	请向管理员申请 namespace get 权限
403	CCI.01.403109	insufficient operation permission, require cci:namespace:list	获取命名空间列表失败, 您可能没有 cci:namespace:list 权限	请向管理员申请 namespace list 权限
403	CCI.01.403110	insufficient operation permission, require cci:namespaceSubResource:create	创建资源失败, 您可能没有 cci:namespaceSubResource:create 权限	请向管理员申请 namespaceSubResource create 权限

状态码	错误码	错误信息	描述	处理措施
403	CCI.01.40311	insufficient operation permission, require cci:namespaceSubResource:delete	删除资源失败, 您可能没有 cci:namespaceSubResource:delete 权限	请向管理员申请 namespaceSubResource delete 权限
403	CCI.01.40312	insufficient operation permission, require cci:namespaceSubResource:update	更新资源失败, 您可能没有 cci:namespaceSubResource:update 权限	请向管理员申请 namespaceSubResource update 权限
403	CCI.01.40313	insufficient operation permission, require cci:namespaceSubResource:get	获取资源失败, 您可能没有 cci:namespaceSubResource:get 权限	请向管理员申请 namespaceSubResource get 权限
403	CCI.01.40314	insufficient operation permission, require cci:namespaceSubResource:list	获取资源列表失败, 您可能没有 cci:namespaceSubResource:list 权限	请向管理员申请 namespaceSubResource list 权限

状态码	错误码	错误信息	描述	处理措施
404	CCI.01.404101	no agency found for user	请求出错：查找的授信关系不存在	请登录前端界面，选到对应region的CCI服务，点击“服务授权”中“同意授权”，CCI将创建委托
401	CCI.01.401102	user has no agency to cci, please apply a agency to cci	请求失败，请先授信CCI服务	请先授信CCI服务
403	CCI.04.403112	Chart must have same name and version with original chart	更新模板失败：模板版本和名称必须与原始模板一致	请确保模板版本和名称必须与原始模板一致
403	CCI.04.403124	The official charts is forbidden to be downloaded	下载模板错误：官方模板禁止下载	请联系技术支持
403	CCI.04.403125	Chart is used, can't deleted	模板删除失败：有应用使用了该模板	请联系技术支持
403	CCI.04.403126	The chart is deprecated, please select a higher version	当前选择的模板版本已废弃，请选择更新的版本	请选择更新的版本

状态码	错误码	错误信息	描述	处理措施
403	CCI.04.403129	Only addon charts have readme.md now, the target chart is not an addon chart	操作失败，readme.md 关联的不是插件模板	请关联插件模板
403	CCI.04.403130	Not allowed to update chart belong to other tenant	请求失败：禁止操作其他用户的模板	禁止操作其他用户的模板
403	CCI.04.403131	Update release is forbidden:The status of release is not DEPLOYED or FAILED	更新失败：当前应用状态不支持更新，安装成功/失败的应用支持更新	请联系技术支持

状态码	错误码	错误信息	描述	处理措施
400	CCI.04.400104	Validate chart version failed	请求失败：模板版本不合法，版本格式要求：长度不大于64位且需满足SemVer语义化版本控制规范。正确示例：1.0.0、1.0.0-alpha、1.0.0-alpha+001	请提供符合版本格式要求的模板版本
400	CCI.04.400105	the ID should not be empty	模板ID不能为空	请输入正确的模板ID
400	CCI.04.400111	Only addon allowed	请求失败：该请求只支持插件模板	请提供插件模板
400	CCI.04.400112	the length of ID is more than 64	请求失败：模板ID长度超过最大64位限制	请输入正确的模板ID
400	CCI.04.400113	Chart ID format is Invalid, must match the regex	请求失败：模板ID格式出错	请输入正确的模板ID
400	CCI.04.400118	Validate chart failed	请求失败：模板文件格式错误，文件后缀名必须是"tgz"，默认最大长度为2兆	请提供文件后缀名是"tgz"，默认长度为2兆内的模板文件

状态码	错误码	错误信息	描述	处理措施
400	CCI.04.400119	Chart file not found	请求失败: 模板文件内容为空	请提供内容不为空的模板文件
400	CCI.04.400121	Chart format is invalid, Chart package name and version must be same with chart name and version	模板包格式错误, 模板包名称和版本必须与模板名称和版本一致	请提供模板包名称和版本与模板名称和版本一致的模板包
400	CCI.04.400126	No target version selected, select one please	请求失败: 未指定插件版本	请指定插件版本
400	CCI.04.400131	Unsupported language	删除readme文件出错: 请求参数的语言模式不正确, 当前仅支持中文和英文	请确保readme文件的语言模式为中文或英文
400	CCI.04.400201	Invalid release name	请求失败: 实例名称不能为空	请输入正确的实例名称
400	CCI.04.400202	Release name length error	请求失败: 实例名长度超过24限制	请输入正确的实例名称

状态码	错误码	错误信息	描述	处理措施
400	CCI.04.400203	Release name not match regex	请求失败：实例名不符合格式要求。格式要求：字母开头，字母/数字结尾，中间支持字母、数字和中划线	请输入正确的实例名称
400	CCI.04.400214	The update action must be either upgrade or rollback	更新应用失败：请求的action不合法，支持：upgrade和rollback	请求的action只支持upgrade和rollback
400	CCI.04.400218	Invalid release version	请求失败：请求的version字段不合法	请输入正确的version字段
400	CCI.04.400302	Failed to Unmarshal	请求失败，请检查body格式是否正确	请提供正确的body格式
400	CCI.04.400304	The request body is too large	请求失败：模板包超过了2M的最大限制	请确保模板包小于2M
400	CCI.04.400306	Validate chart package failed	校验模板包格式时发生错误，请检查模板中yaml文件格式是否正确	请检查模板中yaml文件格式是否正确
404	CCI.04.404204	Release not found	请求失败：查询的应用不存在	请查询正确的应用

状态码	错误码	错误信息	描述	处理措施
409	CCI.04.409112	The chart is already existed in database	模板上传失败，模板名字已存在	请修改模板名字后重试
409	CCI.04.409123	The chart name is conflict with public charts	模板上传失败，模板名与官方模板名冲突	请修改模板名字后重试
409	CCI.04.409213	Release already exists in this cluster	创建实例失败，实例名称重复	请修改实例名称后重试
400	CCI.02.400101	ummarshal ERROR : BadNamespaceRequestBody	创建 namespace 失败：请求结构体错误	根据接口文档排查请求体错误
400	CCI.02.400102	get service account tenant name failed	资源租户获取租户名失败：资源租户不在白名单内	请获取存在白名单内的租户
400	CCI.02.400103	Unsupported Content Type	请求携带文本类型错误	请提供正确的文本类型
400	CCI.02.400104	ummarshal ERROR : BadQuotaRequestBody	创建配额失败：请求结构体错误	请提供正确的结构体

状态码	错误码	错误信息	描述	处理措施
400	CCI.02.400105	update quotas failed	配额更新失败	请联系技术支持
400	CCI.02.400106	used namespaces exceeds quota, could not create any more	创建 namespace 失败: namespace 配额不足	请联系技术支持
400	CCI.02.400107	delete namespace failed, namespace is empty	删除 namespace 失败: namespace 为空, 无法删除	请删除不为空的namespace
400	CCI.02.400108	get flavor info by name failed, flavor name is empty	通过name查询flavor信息失败: name 字段为空	请输入正确的name字段
400	CCI.02.400109	Create namespace failed: NamespaceNameInvalid	创建 namespace 失败: Name字段不合法	请输入正确的name字段
400	CCI.02.400110	Create namespace failed: NamespaceFlavorInvalid	创建 namespace 失败: Flavor字段不合法	请输入正确的Flavor字段

状态码	错误码	错误信息	描述	处理措施
400	CCI.02.400111	Create namespace failed: NamespaceFlavorMissed	创建 namespace 失败: Flavor字段为空	请输入正确的Flavor字段
400	CCI.02.400112	Create namespace failed: NamespaceDomainIDMissed	创建 namespace 失败: DomainID字段为空	请输入正确的DomainID字段
400	CCI.02.400113	Create namespace failed: NamespaceProjectIDMissed	创建 namespace 失败: ProjectID字段为空	请输入正确的ProjectID字段
400	CCI.02.400114	Get namespace failed: NamespaceProjectIDMissed	查询 namespace 失败: ProjectID字段为空	请输入正确的ProjectID字段
400	CCI.02.400116	Get available cluster info from resource manager failed: NoAllocatedCluster	创建 namespace 失败: 当前集群资源不足, 请稍后重试	请联系技术支持

状态码	错误码	错误信息	描述	处理措施
400	CCI.02.400117	Get namespace failed: PaginationParameterInvalid	查询 namespace 失败: 分页参数不合法	请输入正确的分页参数
400	CCI.02.400118	Create namespace failed: EnterpriseProjectIDEmpty	创建 namespace 失败: 企业项目ID为空	请输入正确的企业项目ID
400	CCI.02.400119	Create namespace failed: EnterpriseProjectNotSupported	创建 namespace 失败: 不支持设置企业项目	请联系技术支持
404	CCI.02.404101	flavor not found	查询flavor信息失败: 请求flavor不存在	请查询正确的flavor
404	CCI.02.404102	resourcequota not found	查询资源配额信息失败: 请求资源配额不存在	请查询正确的资源配额
404	CCI.02.404103	quota not found	查询配额信息失败: 请求配额不存在	请查询正确的资源配额
409	CCI.02.409101	namespace name is already exist	创建 namespace 失败: namespace 名称已存在	请修改namespace名称后重试

状态码	错误码	错误信息	描述	处理措施
409	CCI.02.409102	Failed to migrate the namespace	迁移 namespace 失败: namespace 已属于专属节点	请选择非专属命名空间
500	CCI.03.500102	Internal error	内部错误: 与CSB交互时出错	请联系客服或等待一段时间后重试
500	CCI.03.500101	Return error	返回出错: 返回体JSON转换失败	请联系客服或等待一段时间后重试
500	CCI.03.500002	Database request error	数据库请求出错	请联系客服或等待一段时间后重试
500	CCI.03.500001	An internal processing error occurs	内部处理错误	请联系客服或等待一段时间后重试
400	CCI.03.400105	Request error	请求错误: 未携带套餐包信息	请传入有效订单请求
400	CCI.03.400106	Request failed	请求失败: 当前用户无权购买该套餐包	请联系客服申请开通购买权限
400	CCI.03.400107	Request failed	请求失败: 当前用户购买该套餐包总数已达上限	请在当前套餐包使用完后继续购买
400	CCI.03.400108	Request failed	请求失败: 当前用户指定周期内购买该套餐包数量已达上限	请在下个周期内继续购买

状态码	错误码	错误信息	描述	处理措施
400	CCI.02.400122	Failed to migrate the namespace.	迁移 namespace 失败: namespace 状态异常	请修复命名空间状态后重试
400	CCI.02.400121	Failed to migrate the namespace.	迁移 namespace 失败: namespace 下无可用专属节点	请联系客服购买专属节点
400	CCI.03.400109	Request failed	请求失败: 该套餐包剩余可购买数量不足	请扩大套餐包配额或者减少购买量
400	CCI.02.400120	Failed to migrate the namespace.	迁移 namespace 失败: 按需弹性承载参数不合法	请配置有效参数
400	CCI.03.400110	Request failed	请求失败: 无效的活动套餐包	请购买有效套餐包
400	CCI.03.400111	Request failed	请求失败: 活动ID无效, 可能为空或字段过长	请输入有效活动ID
400	CCI.03.400001	The request does not carry a token	请求未携带 token	请在请求体中携带有效token
400	CCI.03.400101	Request error	请求错误: 请求体不合法	请传入有效请求体

状态码	错误码	错误信息	描述	处理措施
400	CCI.03.400102	Request error	请求错误： 请求体JSON 解析失败	请传入有效订单请求
400	CCI.03.400103	Request failed	请求失败： Token缺失用 户信息	请传入有效Token
400	CCI.03.400104	Request failed	请求失败： 无效的POC 套餐包	请购买有效套餐包

8.7 获取项目 ID

操作场景

在调用接口的时候，部分URL中需要填入项目ID，所以需要获取到项目ID。有如下两种获取方式：

- [调用API获取项目ID](#)
- [从控制台获取项目ID](#)

调用 API 获取项目 ID

项目ID可以通过调用[查询指定条件下的项目列表](#)API获取。

获取项目ID的接口为“GET https://{Endpoint}/v3/projects”，其中{Endpoint}为IAM的终端节点，可以从[地区和终端节点](#)获取。接口的认证鉴权请参见[认证鉴权](#)。

响应示例如下，其中projects下的“id”即为项目ID。

```
{
  "projects": [
    {
      "domain_id": "65ewtrgaggshhk1223245sghjlse684b",
      "is_domain": false,
      "parent_id": "65ewtrgaggshhk1223245sghjlse684b",
      "name": "project_name",
      "description": "",
      "links": {
        "next": null,
        "previous": null,
        "self": "https://www.example.com/v3/projects/a4adasfjljaaakla12334jklga9sasfg"
      },
      "id": "a4adasfjljaaakla12334jklga9sasfg",
      "enabled": true
    }
  ],
  "links": {
```

```
"next": null,  
"previous": null,  
"self": "https://www.example.com/v3/projects"  
}  
}
```

从控制台获取项目 ID

从控制台获取项目ID的步骤如下：

1. 登录管理控制台。
2. 鼠标悬停在右上角的用户名，选择下拉列表中的“我的凭证”。
在“API凭证”页面的项目列表中查看项目ID。

图 8-8 查看项目 ID



8.8 获取账号 ID

在调用接口的时候，部分URL中需要填入账号ID（domain-id），所以需要先在管理控制台上获取到账号ID。账号ID获取步骤如下：

1. 登录管理控制台。
2. 单击用户名，在下拉列表中单击“我的凭证”。
在“我的凭证”页面查看账号ID。

图 8-9 获取账号 ID



8.9 获取容器镜像地址

云容器实例支持使用镜像中心和上传到容器镜像服务的镜像。其中

- **容器镜像服务**从镜像中心同步了部分常用镜像，使得您可以在内部网络中直接使用“镜像名称:版本号”，如nginx:alpine，您可以在容器镜像服务控制台中查询同步了哪些镜像。
- 容器镜像服务中的镜像请使用镜像的“下载指令”，上传镜像后，您可以在容器镜像服务的镜像中获取，如下图所示。

图 8-10 镜像地址



9 历史 API

9.1 Kubernetes API（废弃）

📖 说明

当前页面API已经废弃，请使用[Kubernetes API](#)。

9.1.1 Extended PersistentVolumeClaim

9.1.1.1 导入存储

功能介绍

导入已有存储到指定的命名空间。

当前支持导入EVS（云硬盘卷，块存储）、SFS（文件存储卷）、SFS Turbo（极速文件存储卷）、OBS（对象存储卷），使用时 `spec.storageClassName` 参数的取值如下：

- sata：普通I/O云硬盘卷
- sas：高I/O云硬盘卷
- ssd：超高I/O云硬盘卷
- nfs-rw：标准文件协议类型文件存储卷
- efs-performance：性能型极速文件存储卷
- efs-standard：标准型极速文件存储卷
- obs：对象存储卷

📖 说明

- 回收策略采用 **DELETE**，即PVC被删除后，立即删除存储。如果需要保留存储，可使用[解绑存储](#)接口。
- OBS对象存储包含“对象桶”和“并行文件系统”两种类型。并行文件系统是OBS提供了一种经过优化的高性能文件系统，提供毫秒级别访问时延，以及TB/s级别带宽和百万级别的IOPS，相较于OBS对象存储在稳定性、性能上更具优势。因此如需通过挂载方式，**生产环境中推荐您使用OBS并行文件系统**，而不推荐OBS对象存储。

URI

POST /api/v1/namespaces/{namespace}/extended-persistentvolumeclaims

表 9-1 Path 参数

参数	参数类型	描述
namespace	String	命名空间。

请求消息

请求参数

表 9-2 PersistentVolumeClaim v1 数据结构说明

参数	是否必选	参数类型	描述
apiVersion	Yes	String	APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject unrecognized values.
kind	Yes	String	Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to. Cannot be updated.
metadata	Yes	Object	Standard object's metadata. 详情请参见表6-131。 其中annotations字段请参见表9-3。
spec	Yes	Object	Spec defines the desired characteristics of a volume requested by a pod author. 详情请参见表6-298。
status	No	Object	Status represents the current information/status of a persistent volume claim. Read-only. 详情请参见表6-299。

表 9-3 metadata.annotations 字段说明

参数	是否必选	参数类型	描述
fsType	Yes	String	文件系统名称。可设置为： <ul style="list-style-type: none">• 块存储：ext4• 对象存储：obs• 文件存储、极速文件存储：nfs
volumeID	Yes	String	卷id。
deviceMountPath	No	String	共享路径。 说明 只有文件存储卷和极速文件存储卷需要设置该值。

请求示例

```
{
  "apiVersion": "v1",
  "kind": "PersistentVolumeClaim",
  "metadata": {
    "annotations": {
      "fsType": "nfs",
      "volumeID": "378dfa73-3ae4-4179-81c0-67699976b505",
      "deviceMountPath": "sfs-nas01.cn-north-4b.myhuaweicloud.com/share-84b38e4d"
    },
    "name": "cci-sfs-jxre8q80-ylpd",
    "namespace": "test-namespace"
  },
  "spec": {
    "resources": {
      "requests": {
        "storage": "10Gi"
      }
    },
    "storageClassName": "nfs-rw"
  }
}
```

响应消息

响应参数

表 9-4 响应参数

参数	类型	描述
PersistentVolumeClaim	Object	详情请参见 表6-296 。

响应示例

```
{
  "metadata": {
```



```
{
  "name": "pvc-import-efs",
  "namespace": "test-namespace",
  "selfLink": "/api/v1/namespaces/test-namespace/persistentvolumeclaims/pvc-import-efs",
  "uid": "17646a17-a471-11e9-be8a-b44326d0c915",
  "resourceVersion": "65016560",
  "creationTimestamp": "2019-07-12T06:48:44Z",
  "annotations": {
    "kubernetes.io/volumeld": "378dfa73-3ae4-4179-81c0-67699976b505"
  }
},
"spec": {
  "accessModes": [
    "ReadWriteMany"
  ],
  "resources": {
    "requests": {
      "storage": "10"
    }
  },
  "storageClassName": "sata"
},
"status": {
  "phase": "Pending"
}
}
```

状态码

表 9-5 状态码

状态码	描述
200	Export volume success.

更多状态码请参见[状态码](#)。

9.1.1.2 查询导入的 PVC

功能介绍

查询指定命名空间下的PVC。

URI

GET /api/v1/namespaces/{namespace}/extended-persistentvolumeclaims

表 9-6 Path 参数

参数	是否必选	参数类型	描述
namespace	Yes	String	命名空间。

表 9-7 Query 参数

参数	是否必选	参数类型	描述
storageType	No	String	Type of storage, 目前支持: <ul style="list-style-type: none">• bs: 云硬盘存储• nfs: 文件存储• efs: 极速文件存储 如果本参数未指定, 将默认返回当前命名空间下bs(云硬盘存储)类型的PVC信息。

请求消息

N/A

响应消息

响应参数

表 9-8 响应参数

参数	参数类型	描述
PersistentVolumeClaim	Object	详情请参见 表6-296 。
PersistentVolume	Object	详情请参见 表6-305 。
StorageInfo	Object	详情请参见 表6-320 。

响应示例

```
[
  {
    "persistentVolumeClaim": {
      "metadata": {
        "name": "cci-evs-jxzqegxe-k1z3",
        "namespace": "test-namespace",
        "selfLink": "/api/v1/namespaces/csms/persistentvolumeclaims/cci-evs-jxzqegxe-k1z3",
        "uid": "e82c1574-a46d-11e9-be8a-b44326d0c915",
        "resourceVersion": "65003551",
        "creationTimestamp": "2019-07-12T06:25:56Z",
        "annotations": {
          "pv.kubernetes.io/bind-completed": "yes",
          "pv.kubernetes.io/bound-by-controller": "yes",
          "volume.beta.kubernetes.io/storage-provisioner": "flexvolume-huawei.com/fuxivol"
        },
        "finalizers": [
          "kubernetes.io/pvc-protection"
        ],
        "enable": true
      },
      "spec": {
        "accessModes": [
```

```
    "ReadWriteMany"
  ],
  "resources": {
    "requests": {
      "storage": "10Gi"
    }
  },
  "volumeName": "pvc-e82c1574-a46d-11e9-be8a-b44326d0c915",
  "storageClassName": "sata"
},
"status": {
  "phase": "Bound",
  "accessModes": [
    "ReadWriteMany"
  ],
  "capacity": {
    "storage": "10Gi"
  }
}
},
"persistentVolume": {
  "metadata": {
    "name": "pvc-e82c1574-a46d-11e9-be8a-b44326d0c915",
    "selfLink": "/api/v1/persistentvolumes/pvc-e82c1574-a46d-11e9-be8a-b44326d0c915",
    "uid": "eb0a0ca1-a46d-11e9-be8a-b44326d0c915",
    "resourceVersion": "65003549",
    "creationTimestamp": "2019-07-12T06:26:01Z",
    "labels": {
      "tenant.kubernetes.io/domain-id": "f0c61dbd65974140956ed37a91ea860f",
      "tenant.kubernetes.io/project-id": "cdb4249297a44665a63eec4f27ad09bf"
    },
  },
  "annotations": {
    "kubernetes.io/createdby": "huawei.com/fuxivol-dynamic-provisioner",
    "pv.kubernetes.io/bound-by-controller": "yes",
    "pv.kubernetes.io/namespace": "test-namespace",
    "pv.kubernetes.io/provisioned-by": "flexvolume-huawei.com/fuxivol",
    "tenant.kubernetes.io/domain-id": "f0c61dbd65974140956ed37a91ea860f",
    "tenant.kubernetes.io/project-id": "cdb4249297a44665a63eec4f27ad09bf"
  },
  "finalizers": [
    "kubernetes.io/pv-protection"
  ]
},
"spec": {
  "capacity": {
    "storage": "10Gi"
  },
  "flexVolume": {
    "driver": "huawei.com/fuxivol",
    "fsType": "ext4",
    "options": {
      "fsType": "ext4",
      "volumeID": "06e10708-6412-4190-8496-f9531fb5fd0c"
    }
  },
  "accessModes": [
    "ReadWriteMany"
  ],
  "claimRef": {
    "kind": "PersistentVolumeClaim",
    "namespace": "test-namespace",
    "name": "cci-evs-jxzqegxe-k1z3",
    "uid": "e82c1574-a46d-11e9-be8a-b44326d0c915",
    "apiVersion": "v1",
    "resourceVersion": "65003516"
  },
  "persistentVolumeReclaimPolicy": "Delete",
  "storageClassName": "sata"
},
}
```

```
    "status": {
      "phase": "Bound"
    }
  },
  "storageInfo": {
    "kind": "Volume",
    "apiVersion": "paas/v1beta1",
    "metadata": {
      "name": "pvc-e82c1574-a46d-11e9-be8a-b44326d0c915",
      "creationTimestamp": "2019-07-12T06:25:56Z",
      "labels": {
        "_system_volume_name": "pvc-e82c1574-a46d-11e9-be8a-b44326d0c915",
        "hw.passthrough": "true",
        "kubernetes.io/namespace": "test-namespace",
        "tenant.kubernetes.io/domain-id": "f0c61dbd65974140956ed37a91ea860f",
        "tenant.kubernetes.io/project-id": "cdb4249297a44665a63eec4f27ad09bf"
      }
    }
  },
  "spec": {
    "name": "pvc-e82c1574-a46d-11e9-be8a-b44326d0c915",
    "size": 10,
    "inresourcepool": false,
    "availability_zone": "cn-north-1a",
    "volume_type": "sata",
    "multiattach": true,
    "access": {
      "": {
        "access_type": "",
        "access_to": "",
        "access_level": "",
        "id": "",
        "state": ""
      }
    }
  },
  "vpc_id": ""
},
"status": {
  "id": "06e10708-6412-4190-8496-f9531fb5fd0c",
  "status": "available",
  "created_at": "2019-07-12T06:25:56Z",
  "attachments": null,
  "app_info": null
}
}
]
```

状态码

表 9-9 状态码

状态码	描述
200	查询成功，并返回结果。

更多状态码请参见[状态码](#)。

9.1.1.3 解绑存储

功能描述

从指定命名空间解绑存储的接口。

URI

DELETE /api/v1/namespaces/{namespace}/persistentvolumeclaims/{name}

表 9-10 Path 参数

参数	是否必选	参数类型	描述
namespace	Yes	String	命名空间。
name	Yes	String	persistentvolumeclaims名称

表 9-11 Query 参数

参数	是否必选	参数类型	描述
deleteVolume	No	Boolean	是否删除volume，可选值为true和false： <ul style="list-style-type: none">• true：同时删除使用的存储• false：保留存储，只删除pvc
storageType	No	String	Type of storage，目前支持： <ul style="list-style-type: none">• bs: 云硬盘存储• nfs: 文件存储 说明 当deleteVolume=true时，storageType为非必填项；当deleteVolume=false时，storageType为必填项。

请求消息

N/A

响应消息

N/A

状态码

表 9-12 状态码

状态码	描述
200	删除成功

更多状态码请参见[状态码](#)。

9.1.2 TFJob

9.1.2.1 创建 TFJob

功能介绍

创建TFJob。

TFJob即Tensorflow任务，是基于Tensorflow开源框架的kubernetes自定义资源类型，有多种角色可以配置，能够帮助我们更简单地实现Tensorflow的单机或分布式训练。Tensorflow开源框架的信息详见：<https://www.tensorflow.org>。

URI

POST /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs

表 9-13 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-14 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数：

请求参数的详细描述请参见[表6-275](#)。

请求示例：

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "TFJob",
  "metadata": {
    "name": "tfjob-test"
  },
  "spec": {
    "backoffLimit": 6,
    "tfReplicaSpecs": {
      "Ps": {
        "replicas": 1,
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "python",
```

```
        "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
        "--batch_size=1",
        "--model=resnet50",
        "--variable_update=parameter_server",
        "--flush_stdout=true",
        "--num_gpus=1",
        "--local_parameter_device=cpu",
        "--device=cpu",
        "--data_format=NHWC"
    ],
    "image": "k8s.gcr.io/tf-benchmarks-cpu:v1",
    "name": "tensorflow",
    "ports": [
        {
            "containerPort": 2222,
            "name": "tfjob-port"
        }
    ],
    "resources": {
        "limits": {
            "cpu": "2",
            "memory": "4Gi"
        },
        "requests": {
            "cpu": "2",
            "memory": "4Gi"
        }
    }
},
"restartPolicy": "OnFailure",
"imagePullSecrets": [
    {
        "name": "imagepull-secret"
    }
]
}
},
"Worker": {
    "replicas": 1,
    "template": {
        "spec": {
            "containers": [
                {
                    "args": [
                        "python",
                        "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                        "--batch_size=1",
                        "--model=resnet50",
                        "--variable_update=parameter_server",
                        "--flush_stdout=true",
                        "--local_parameter_device=cpu",
                        "--device=cpu",
                        "--data_format=NHWC"
                    ],
                    "image": "k8s.gcr.io/tf-benchmarks-cpu:v1",
                    "name": "tensorflow",
                    "ports": [
                        {
                            "containerPort": 2222,
                            "name": "tfjob-port"
                        }
                    ],
                    "resources": {
                        "limits": {
                            "cpu": "2",
                            "memory": "4Gi"
                        },
```

```
        "requests": {
            "cpu": "2",
            "memory": "4Gi"
        }
    },
    ],
    "restartPolicy": "OnFailure",
    "imagePullSecrets": [
        {
            "name": "imagepull-secret"
        }
    ]
}
}
}
}
}
}
```

响应消息

响应参数:

响应参数的详细描述请参考[表6-275](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "TFJob",
  "metadata": {
    "creationTimestamp": "2019-07-23T12:39:47Z",
    "generation": 1,
    "name": "tfjob-test",
    "namespace": "kube-test",
    "resourceVersion": "72050567",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",
    "uid": "f461f966-ad46-11e9-aaa4-340a9837e413"
  },
  "spec": {
    "backoffLimit": 6,
    "tfReplicaSpecs": {
      "Ps": {
        "replicas": 1,
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "python",
                  "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                  "--batch_size=1",
                  "--model=resnet50",
                  "--variable_update=parameter_server",
                  "--flush_stdout=true",
                  "--num_gpus=1",
                  "--local_parameter_device=cpu",
                  "--device=cpu",
                  "--data_format=NHWC"
                ],
                "image": "gcr.io/ml-org:20190723/tf-benchmarks-cpu:v1",
                "name": "tensorflow",
                "ports": [
                  {
                    "containerPort": 2222,
                    "name": "tfjob-port"
                  }
                ]
              }
            ],
            "restartPolicy": "OnFailure"
          }
        }
      }
    }
  }
}
```



```
"status": {  
  }  
}
```

状态码

表 9-15 状态码

状态码	描述
200	OK
201	Created
202	Accepted
401	Unauthorized
400	Badrequest
500	Internal error
403	Forbidden

9.1.2.2 查询 TFJob

功能介绍

查询TFJob的详细信息。

URI

GET /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs/{name}

表 9-16 Path 参数

参数	是否必选	描述
name	Yes	name of the TFJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-17 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参考[表6-275](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "TFJob",
  "metadata": {
    "creationTimestamp": "2019-07-23T12:39:47Z",
    "generation": 1,
    "name": "tfjob-test",
    "namespace": "kube-test",
    "resourceVersion": "72050567",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",
    "uid": "f461f966-ad46-11e9-aaa4-340a9837e413"
  },
  "spec": {
    "backoffLimit": 6,
    "tfReplicaSpecs": {
      "Ps": {
        "replicas": 1,
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "python",
                  "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                  "--batch_size=1",
                  "--model=resnet50",
                  "--variable_update=parameter_server",
                  "--flush_stdout=true",
                  "--num_gpus=1",
                  "--local_parameter_device=cpu",
                  "--device=cpu",
                  "--data_format=NHWC"
                ],
                "image": "**.215:20202/ci/tf-benchmarks-cpu:v1",
                "name": "tensorflow",
                "ports": [
                  {
                    "containerPort": 2222,
                    "name": "tfjob-port"
                  }
                ],
                "resources": {
                  "limits": {
                    "cpu": "2",
                    "memory": "4Gi"
                  },
                  "requests": {
                    "cpu": "2",
                    "memory": "4Gi"
                  }
                }
              }
            ]
          }
        }
      }
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
}
```

```
    }
  ],
  "restartPolicy": "OnFailure"
}
},
"Worker": {
  "replicas": 1,
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "python",
            "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
            "--batch_size=1",
            "--model=resnet50",
            "--variable_update=parameter_server",
            "--flush_stdout=true",
            "--local_parameter_device=cpu",
            "--device=cpu",
            "--data_format=NHWC"
          ],
          "image": "***215:20202/ci/tf-benchmarks-cpu:v1",
          "name": "tensorflow",
          "ports": [
            {
              "containerPort": 2222,
              "name": "tfjob-port"
            }
          ],
          "resources": {
            "limits": {
              "cpu": "2",
              "memory": "4Gi"
            },
            "requests": {
              "cpu": "2",
              "memory": "4Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ],
    "restartPolicy": "OnFailure"
  }
}
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-23T12:38:58Z",
      "lastUpdateTime": "2019-07-23T12:38:58Z",
      "message": "TFJob tfjob-test is created.",
      "reason": "TFJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-23T12:39:30Z",
      "lastUpdateTime": "2019-07-23T12:39:30Z",
      "message": "TFJob tfjob-test is running.",
      "reason": "TFJobRunning",
```

```

        "status": "True",
        "type": "Running"
    }
  ],
  "replicaStatuses": {
    "PS": {
      "active": 1
    },
    "Worker": {
      "active": 1
    }
  },
  "startTime": "2019-07-23T12:38:58Z"
}
}

```

状态码

表 9-18 状态码

状态码	描述
200	OK
401	Unauthorized
404	Not found
500	Internal error
403	Forbidden

9.1.2.3 查询指定 namespace 下的所有 TFJob

功能介绍

查询Namespace下所有TFJob的详细信息。

URI

GET /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs

表 9-19 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-20 Query 参数

参数	是否必选	描述
fieldSelector	No	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	No	A selector to restrict the list of returned objects by their labels. Defaults to everything.
limit	No	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the continue field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
resourceVersion	No	When specified with a watch call, shows changes that occur after that particular version of a resource. Defaults to changes from the beginning of history. When specified for list: - if unset, then the result is returned from remote storage based on quorum-read flag; - if it's 0, then we simply return what we currently have in cache, no guarantee; - if set to non zero, then the result is at least as fresh as given rv.
timeoutSeconds	No	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

参数	是否必选	描述
watch	No	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-289](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "items": [
    {
      "apiVersion": "kubeflow.org/v1",
      "kind": "TFJob",
      "metadata": {
        "creationTimestamp": "2019-07-23T12:39:47Z",
        "generation": 1,
        "name": "tfjob-test",
        "namespace": "kube-test",
        "resourceVersion": "72050567",
        "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",
        "uid": "f461f966-ad46-11e9-aaa4-340a9837e413"
      },
      "spec": {
        "backoffLimit": 6,
        "tfReplicaSpecs": {
          "Ps": {
            "replicas": 1,
            "template": {
              "spec": {
                "containers": [
                  {
                    "args": [
                      "python",
                      "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                      "--batch_size=1",
                      "--model=resnet50",
                      "--variable_update=parameter_server",
                      "--flush_stdout=true",
                      "--num_gpus=1",
                      "--local_parameter_device=cpu",
                      "--device=cpu",
                      "--data_format=NHWC"
                    ],
                    "image": "*/.*/215:20202/cci/tf-benchmarks-cpu:v1",
                    "name": "tensorflow",
                    "ports": [
                      {
                        "containerPort": 2222,
                        "name": "tfjob-port"
                      }
                    ]
                  }
                ]
              },
              "resources": {
```

```
        "limits": {
            "cpu": "2",
            "memory": "4Gi"
        },
        "requests": {
            "cpu": "2",
            "memory": "4Gi"
        }
    }
},
"imagePullSecrets": [
    {
        "name": "imagepull-secret"
    }
],
"restartPolicy": "OnFailure"
}
},
"Worker": {
    "replicas": 1,
    "template": {
        "spec": {
            "containers": [
                {
                    "args": [
                        "python",
                        "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                        "--batch_size=1",
                        "--model=resnet50",
                        "--variable_update=parameter_server",
                        "--flush_stdout=true",
                        "--local_parameter_device=cpu",
                        "--device=cpu",
                        "--data_format=NHWC"
                    ],
                    "image": "**.*.215:20202/ccl/tf-benchmarks-cpu:v1",
                    "name": "tensorflow",
                    "ports": [
                        {
                            "containerPort": 2222,
                            "name": "tfjob-port"
                        }
                    ],
                    "resources": {
                        "limits": {
                            "cpu": "2",
                            "memory": "4Gi"
                        },
                        "requests": {
                            "cpu": "2",
                            "memory": "4Gi"
                        }
                    }
                }
            ],
            "imagePullSecrets": [
                {
                    "name": "imagepull-secret"
                }
            ],
            "restartPolicy": "OnFailure"
        }
    }
},
"status": {
```



```
    "conditions": [
      {
        "lastTransitionTime": "2019-07-23T12:38:58Z",
        "lastUpdateTime": "2019-07-23T12:38:58Z",
        "message": "TFJob tfjob-test is created.",
        "reason": "TFJobCreated",
        "status": "True",
        "type": "Created"
      },
      {
        "lastTransitionTime": "2019-07-23T12:39:30Z",
        "lastUpdateTime": "2019-07-23T12:39:30Z",
        "message": "TFJob tfjob-test is running.",
        "reason": "TFJobRunning",
        "status": "True",
        "type": "Running"
      }
    ],
    "replicaStatuses": {
      "PS": {
        "active": 1
      },
      "Worker": {
        "active": 1
      }
    },
    "startTime": "2019-07-23T12:38:58Z"
  }
},
"kind": "TFJobList",
"metadata": {
  "continue": "",
  "resourceVersion": "72353810",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs"
}
}
```

状态码

表 9-21 状态码

状态码	描述
200	OK
401	Unauthorized
404	Not found
500	Internal error

9.1.2.4 删除 namespace 下的所有 TFJob

功能介绍

删除命名空间下的所有 TFJob。

URI

DELETE /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs

表 9-22 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-23 Query 参数

参数	是否必选	描述
fieldSelector	No	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	No	A selector to restrict the list of returned objects by their labels. Defaults to everything.
limit	No	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the continue field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>

参数	是否必选	描述
resourceVersion	No	When specified with a watch call, shows changes that occur after that particular version of a resource. Defaults to changes from the beginning of history. When specified for list: - if unset, then the result is returned from remote storage based on quorum-read flag; - if it's 0, then we simply return what we currently have in cache, no guarantee; - if set to non zero, then the result is at least as fresh as given rv.
timeoutSeconds	No	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.
watch	No	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求消息

请求参数:

N/A

响应消息

响应参数:

响应参数的详细描述请参考[表6-289](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "items": [
    {
      "apiVersion": "kubeflow.org/v1",
      "kind": "TFJob",
      "metadata": {
        "creationTimestamp": "2019-07-23T12:39:47Z",
        "generation": 1,
        "name": "tfjob-test",
        "namespace": "kube-test",
        "resourceVersion": "72050567",
        "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",
        "uid": "f461f966-ad46-11e9-aaa4-340a9837e413"
      },
      "spec": {
        "backoffLimit": 6,
        "tfReplicaSpecs": {
          "Ps": {
            "replicas": 1,
            "template": {
              "spec": {
                "containers": [
                  {
                    "args": [
```

```
        "python",
        "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
        "--batch_size=1",
        "--model=resnet50",
        "--variable_update=parameter_server",
        "--flush_stdout=true",
        "--num_gpus=1",
        "--local_parameter_device=cpu",
        "--device=cpu",
        "--data_format=NHWC"
    ],
    "image": "**:*:215:20202/cci/tf-benchmarks-cpu:v1",
    "name": "tensorflow",
    "ports": [
        {
            "containerPort": 2222,
            "name": "tfjob-port"
        }
    ],
    "resources": {
        "limits": {
            "cpu": "2",
            "memory": "4Gi"
        },
        "requests": {
            "cpu": "2",
            "memory": "4Gi"
        }
    }
},
"imagePullSecrets": [
    {
        "name": "imagepull-secret"
    }
],
"restartPolicy": "OnFailure"
}
},
"Worker": {
    "replicas": 1,
    "template": {
        "spec": {
            "containers": [
                {
                    "args": [
                        "python",
                        "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                        "--batch_size=1",
                        "--model=resnet50",
                        "--variable_update=parameter_server",
                        "--flush_stdout=true",
                        "--local_parameter_device=cpu",
                        "--device=cpu",
                        "--data_format=NHWC"
                    ],
                    "image": "**:*:215:20202/cci/tf-benchmarks-cpu:v1",
                    "name": "tensorflow",
                    "ports": [
                        {
                            "containerPort": 2222,
                            "name": "tfjob-port"
                        }
                    ],
                    "resources": {
                        "limits": {
                            "cpu": "2",
                            "memory": "4Gi"
                        }
                    }
                }
            ]
        }
    }
}
```

```
    },
    "requests": {
      "cpu": "2",
      "memory": "4Gi"
    }
  }
},
"imagePullSecrets": [
  {
    "name": "imagepull-secret"
  }
],
"restartPolicy": "OnFailure"
}
}
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-23T12:38:58Z",
      "lastUpdateTime": "2019-07-23T12:38:58Z",
      "message": "TFJob tfjob-test is created.",
      "reason": "TFJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-23T12:39:30Z",
      "lastUpdateTime": "2019-07-23T12:39:30Z",
      "message": "TFJob tfjob-test is running.",
      "reason": "TFJobRunning",
      "status": "True",
      "type": "Running"
    }
  ],
  "replicaStatuses": {
    "PS": {
      "active": 1
    },
    "Worker": {
      "active": 1
    }
  },
  "startTime": "2019-07-23T12:38:58Z"
}
},
"kind": "TFJobList",
"metadata": {
  "continue": "",
  "resourceVersion": "72353810",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs"
}
}
```

状态码

表 9-24 状态码

状态码	描述
200	OK

状态码	描述
401	Unauthorized
500	Internal error

9.1.2.5 删除 TFJob

功能介绍

删除TFJob。

URI

DELETE /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs/{name}

表 9-25 Path 参数

参数	是否必选	描述
name	Yes	name of the TFJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-26 Query 参数

参数	是否必选	描述
dryRun	No	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	No	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	No	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	描述
propagationPolicy	No	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	No	If 'true' , then the output is pretty printed.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-193](#)。

响应示例:

```
{
  "kind": "Status",
  "apiVersion": "v1",
  "metadata": {},
  "status": "Success",
  "details": {
    "name": "tfjob-test",
    "group": "kubeflow.org",
    "kind": "tfjobs",
    "uid": "f461f966-ad46-11e9-aaa4-340a9837e413"
  }
}
```

状态码

表 9-27 状态码

状态码	描述
200	OK
202	Accepted
401	Unauthorized
500	Internal Error
403	Forbidden

9.1.2.6 更新 TFJob

功能介绍

更新TFJob。如下字段可被更新：

- metadata.labels
- metadata.annotations
- spec.activeDeadlineSeconds
- spec.ttlSecondsAfterFinished
- spec.cleanPodPolicy

URI

PATCH /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs/{name}

表 9-28 Path 参数

参数	是否必选	描述
name	Yes	name of the TFJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-29 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数：

“Content-Type” 消息头说明请参见[PATCH请求方法操作说明](#)。

📖 说明

目前只支持“Merge Patch”。

请求示例：

```
Content-Type: application/merge-patch+json
{
  "metadata": {
    "labels": {
      "app": "test"
    }
  }
}
```



```
}  
}
```

响应消息

响应参数:

响应参数的详细描述请参见[表6-275](#)。

响应示例:

```
{  
  "apiVersion": "kubeflow.org/v1",  
  "kind": "TFJob",  
  "metadata": {  
    "creationTimestamp": "2019-07-24T07:17:01.000Z",  
    "generation": 1,  
    "labels": {  
      "app": "test"  
    },  
    "name": "tfjob-test",  
    "namespace": "kube-test",  
    "resourceVersion": "72444814",  
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",  
    "uid": "083cc6df-ade3-11e9-aaa4-340a9837e413"  
  },  
  "spec": {  
    "backoffLimit": 6,  
    "tfReplicaSpecs": {  
      "Ps": {  
        "replicas": 1,  
        "template": {  
          "spec": {  
            "containers": [  
              {  
                "args": [  
                  "python",  
                  "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",  
                  "--batch_size=1",  
                  "--model=resnet50",  
                  "--variable_update=parameter_server",  
                  "--flush_stdout=true",  
                  "--num_gpus=1",  
                  "--local_parameter_device=cpu",  
                  "--device=cpu",  
                  "--data_format=NHWC"  
                ],  
                "image": "*/*/215:20202/cci/tf-benchmarks-cpu:v1",  
                "name": "tensorflow",  
                "ports": [  
                  {  
                    "containerPort": 2222,  
                    "name": "tfjob-port"  
                  }  
                ],  
                "resources": {  
                  "limits": {  
                    "cpu": "2",  
                    "memory": "4Gi"  
                  },  
                  "requests": {  
                    "cpu": "2",  
                    "memory": "4Gi"  
                  }  
                }  
              }  
            ]  
          }  
        }  
      }  
    },  
    "imagePullSecrets": [  
      {  
        "name": "secret-name"  
      }  
    ]  
  }  
}
```

```
        "name": "imagepull-secret"
      }
    ],
    "restartPolicy": "OnFailure"
  }
},
"Worker": {
  "replicas": 1,
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "python",
            "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
            "--batch_size=1",
            "--model=resnet50",
            "--variable_update=parameter_server",
            "--flush_stdout=true",
            "--local_parameter_device=cpu",
            "--device=cpu",
            "--data_format=NHWC"
          ],
          "image": "**.215:20202/ci/tf-benchmarks-cpu:v1",
          "name": "tensorflow",
          "ports": [
            {
              "containerPort": 2222,
              "name": "tfjob-port"
            }
          ],
          "resources": {
            "limits": {
              "cpu": "2",
              "memory": "4Gi"
            },
            "requests": {
              "cpu": "2",
              "memory": "4Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ],
    "restartPolicy": "OnFailure"
  }
}
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T07:16:13.000Z",
      "lastUpdateTime": "2019-07-24T07:16:13.000Z",
      "message": "TFJob tfjob-test is created.",
      "reason": "TFJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-24T07:16:18.000Z",
      "lastUpdateTime": "2019-07-24T07:16:18.000Z",
      "message": "TFJob tfjob-test is running."
    }
  ]
}
```

```
    "reason": "TFJobRunning",  
    "status": "True",  
    "type": "Running"  
  }  
],  
"replicaStatuses": {  
  "PS": {  
    "active": 1  
  },  
  "Worker": {  
    "active": 1  
  }  
},  
"startTime": "2019-07-24T07:16:13.000Z"  
}
```

状态码

表 9-30 状态码

状态码	描述
200	OK
401	Unauthorized
500	Internal Error
403	Forbidden
409	Conflict
400	BadRequest

9.1.2.7 替换 TFJob

功能介绍

替换TFJob。如下字段可被替换：

- metadata.labels
- metadata.annotations
- spec.activeDeadlineSeconds
- spec.ttlSecondsAfterFinished
- spec.cleanPodPolicy

URI

PUT /apis/kubeflow.org/v1/namespaces/{namespace}/tfjobs/{name}

表 9-31 Path 参数

参数	是否必选	描述
name	Yes	name of the TFJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-32 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数的详细描述请参见[表6-275](#)。

请求示例：

更改TFJob的结束存活时间ttlSecondsAfterFinished：

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "TFJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T07:17:01Z",
    "generation": 2,
    "labels": {
      "app": "test"
    },
    "name": "tfjob-test",
    "namespace": "kube-test",
    "resourceVersion": "72447176",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",
    "uid": "083cc6df-ade3-11e9-aaa4-340a9837e413"
  },
  "spec": {
    "backoffLimit": 6,
    "tfReplicaSpecs": {
      "Ps": {
        "replicas": 1,
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "python",
                  "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
                  "--batch_size=1",
                  "--model=resnet50",
                  "--variable_update=parameter_server",
                  "--flush_stdout=true",
                  "--num_gpus=1",
                  "--local_parameter_device=cpu",
                  "--device=cpu",
                  "--data_format=NHWC"
                ],
                "image": "**.*.215:20202/ci/tf-benchmarks-cpu:v1",
                "name": "tensorflow",
```

```
        "ports": [
          {
            "containerPort": 2222,
            "name": "tfjob-port"
          }
        ],
        "resources": {
          "limits": {
            "cpu": "2",
            "memory": "4Gi"
          },
          "requests": {
            "cpu": "2",
            "memory": "4Gi"
          }
        }
      }
    ],
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ],
    "restartPolicy": "OnFailure"
  }
},
"Worker": {
  "replicas": 1,
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "python",
            "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
            "--batch_size=1",
            "--model=resnet50",
            "--variable_update=parameter_server",
            "--flush_stdout=true",
            "--local_parameter_device=cpu",
            "--device=cpu",
            "--data_format=NHWC"
          ],
          "image": "k8s.gcr.io/tf-benchmarks-cpu:v1",
          "name": "tensorflow",
          "ports": [
            {
              "containerPort": 2222,
              "name": "tfjob-port"
            }
          ],
          "resources": {
            "limits": {
              "cpu": "2",
              "memory": "4Gi"
            },
            "requests": {
              "cpu": "2",
              "memory": "4Gi"
            }
          }
        }
      ]
    }
  },
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ],
}
```

```
        "restartPolicy": "OnFailure"
      }
    }
  },
  "ttlSecondsAfterFinished": 1000
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T07:16:13Z",
      "lastUpdateTime": "2019-07-24T07:16:13Z",
      "message": "TFJob tfjob-test is created.",
      "reason": "TFJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-24T07:16:18Z",
      "lastUpdateTime": "2019-07-24T07:16:18Z",
      "message": "TFJob tfjob-test is running.",
      "reason": "TFJobRunning",
      "status": "True",
      "type": "Running"
    }
  ],
  "replicaStatuses": {
    "PS": {
      "active": 1
    },
    "Worker": {
      "active": 1
    }
  },
  "startTime": "2019-07-24T07:16:13Z"
}
}
```

响应消息

响应参数:

响应参数的详细描述请参考[表6-275](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "TFJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T07:17:01Z",
    "generation": 2,
    "labels": {
      "app": "test"
    },
    "name": "tfjob-test",
    "namespace": "kube-test",
    "resourceVersion": "72447176",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/tfjobs/tfjob-test",
    "uid": "083cc6df-ade3-11e9-aaa4-340a9837e413"
  },
  "spec": {
    "backoffLimit": 6,
    "tfReplicaSpecs": {
      "Ps": {
        "replicas": 1,
        "template": {
          "spec": {
            "containers": [
```

```
{
  "args": [
    "python",
    "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
    "--batch_size=1",
    "--model=resnet50",
    "--variable_update=parameter_server",
    "--flush_stdout=true",
    "--num_gpus=1",
    "--local_parameter_device=cpu",
    "--device=cpu",
    "--data_format=NHWC"
  ],
  "image": "**.*.215:20202/cqi/tf-benchmarks-cpu:v1",
  "name": "tensorflow",
  "ports": [
    {
      "containerPort": 2222,
      "name": "tfjob-port"
    }
  ],
  "resources": {
    "limits": {
      "cpu": "2",
      "memory": "4Gi"
    },
    "requests": {
      "cpu": "2",
      "memory": "4Gi"
    }
  }
},
"imagePullSecrets": [
  {
    "name": "imagepull-secret"
  }
],
"restartPolicy": "OnFailure"
}
},
"Worker": {
  "replicas": 1,
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "python",
            "/opt/tf-benchmarks/scripts/tf_cnn_benchmarks/tf_cnn_benchmarks.py",
            "--batch_size=1",
            "--model=resnet50",
            "--variable_update=parameter_server",
            "--flush_stdout=true",
            "--local_parameter_device=cpu",
            "--device=cpu",
            "--data_format=NHWC"
          ],
          "image": "**.*.215:20202/cqi/tf-benchmarks-cpu:v1",
          "name": "tensorflow",
          "ports": [
            {
              "containerPort": 2222,
              "name": "tfjob-port"
            }
          ],
          "resources": {
            "limits": {
```


状态码	描述
500	Internal Error
403	Forbidden

9.1.3 MXJob

9.1.3.1 创建 MXJob

功能介绍

创建MXJob。

MXJob即MXNet任务，是基于MXNet开源框架的kubernetes自定义资源类型，有多种角色可以配置，能够帮助我们更简单地实现MXNet的训练。MXNet开源框架的信息详见：<https://mxnet.incubator.apache.org/>。

URI

POST /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs

表 9-34 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-35 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数：

请求参数的详细描述请参考[表6-283](#)。

请求示例：

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "MXJob",
  "metadata": {
    "name": "mxnet-job"
  },
}
```

```
"spec": {
  "cleanPodPolicy": "Running",
  "jobMode": "MXTrain",
  "mxReplicaSpecs": {
    "Scheduler": {
      "replicas": 1,
      "restartPolicy": "Never",
      "template": {
        "spec": {
          "imagePullSecrets": [
            {
              "name": "imagepull-secret"
            }
          ],
          "containers": [
            {
              "name": "mxnet",
              "image": "**:*:215:20202/cci/mxnet:xsw-dis",
              "command": [
                "/bin/bash"
              ],
              "args": [
                "-c",
                "python train_imagenet.py"
              ],
              "resources": {
                "requests": {
                  "cpu": "1000m",
                  "memory": "2Gi"
                },
                "limits": {
                  "cpu": "1000m",
                  "memory": "2Gi"
                }
              }
            }
          ]
        }
      }
    }
  }
},
"Server": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ],
      "containers": [
        {
          "name": "mxnet",
          "image": "**:*:215:20202/cci/mxnet:xsw-dis",
          "command": [
            "/bin/bash"
          ],
          "args": [
            "-c",
            "python train_imagenet.py"
          ],
          "resources": {
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ]
    }
  }
}
```

```
    }
  }
}
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ],
      "containers": [
        {
          "name": "mxnet",
          "image": "**.*.215:20202/cci/mxnet:xsw-dis",
          "command": [
            "/bin/bash"
          ],
          "args": [
            "-c",
            "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
          ],
          "resources": {
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ]
    }
  }
}
}
```

响应消息

响应参数:

响应参数的详细描述请参考[表6-283](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "MXJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T08:42:33Z",
    "generation": 1,
    "name": "mxnet-job",
    "namespace": "kube-test",
    "resourceVersion": "72476154",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",
    "uid": "fac6dcd2-adee-11e9-8041-340a9837e2a7"
  },
  "spec": {
```

```
"cleanPodPolicy": "Running",
"jobMode": "MXTrain",
"mxReplicaSpecs": {
  "Scheduler": {
    "replicas": 1,
    "restartPolicy": "Never",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "-c",
              "python train_imagenet.py"
            ],
            "command": [
              "/bin/bash"
            ],
            "image": "***.215:20202/cqi/mxnet:xsw-dis",
            "name": "mxnet",
            "resources": {
              "limits": {
                "cpu": "1000m",
                "memory": "2Gi"
              },
              "requests": {
                "cpu": "1000m",
                "memory": "2Gi"
              }
            }
          }
        ]
      }
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"Server": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "***.215:20202/cqi/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
```

```

        "name": "imagepull-secret"
      }
    ]
  },
  "Worker": {
    "replicas": 1,
    "restartPolicy": "Never",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "-c",
              "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
            ],
            "command": [
              "/bin/bash"
            ],
            "image": "**.*.215:20202/cci/mxnet:xsw-dis",
            "name": "mxnet",
            "resources": {
              "limits": {
                "cpu": "1000m",
                "memory": "2Gi"
              },
              "requests": {
                "cpu": "1000m",
                "memory": "2Gi"
              }
            }
          }
        ]
      }
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"status": {
}
}

```

状态码

表 9-36 状态码

状态码	描述
200	OK
201	Created
202	Accepted
401	Unauthorized
400	Badrequest

状态码	描述
500	Internal error
403	Forbidden

9.1.3.2 查询 MXJob

功能介绍

查询MXJob的详细信息。

URI

GET /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs/{name}

表 9-37 Path 参数

参数	是否必选	描述
name	Yes	name of the MXJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-38 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参考[表6-283](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "MXJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T08:42:33Z",
    "generation": 1,
    "name": "mxnet-job",
    "namespace": "kube-test",
```

```
"resourceVersion": "72476154",
"selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",
"uid": "fac6dcd2-adee-11e9-8041-340a9837e2a7"
},
"spec": {
  "cleanPodPolicy": "Running",
  "jobMode": "MXTrain",
  "mxReplicaSpecs": {
    "Scheduler": {
      "replicas": 1,
      "restartPolicy": "Never",
      "template": {
        "spec": {
          "containers": [
            {
              "args": [
                "-c",
                "python train_imagenet.py"
              ],
              "command": [
                "/bin/bash"
              ],
              "image": "*/*.215:20202/cci/mxnet:xsw-dis",
              "name": "mxnet",
              "resources": {
                "limits": {
                  "cpu": "1000m",
                  "memory": "2Gi"
                },
                "requests": {
                  "cpu": "1000m",
                  "memory": "2Gi"
                }
              }
            }
          ]
        },
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      }
    }
  },
  "Server": {
    "replicas": 1,
    "restartPolicy": "Never",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "-c",
              "python train_imagenet.py"
            ],
            "command": [
              "/bin/bash"
            ],
            "image": "*/*.215:20202/cci/mxnet:xsw-dis",
            "name": "mxnet",
            "resources": {
              "limits": {
                "cpu": "1000m",
                "memory": "2Gi"
              },
              "requests": {
                "cpu": "1000m",
                "memory": "2Gi"
              }
            }
          }
        ]
      }
    }
  }
}
```

```
    }
  },
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ]
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "*/*.215:20202/ci/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ]
    }
  },
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ]
}
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T08:42:33Z",
      "lastUpdateTime": "2019-07-24T08:42:33Z",
      "message": "MXJob mxnet-job is created.",
      "reason": "MXJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-24T08:42:33Z",
      "lastUpdateTime": "2019-07-24T08:42:52Z",
      "message": "MXJob mxnet-job is running.",
      "reason": "MXJobRunning",
      "status": "True",
      "type": "Running"
    }
  ]
},
],
```



```
"mxReplicaStatuses": {  
  "Scheduler": {  
    "active": 1  
  },  
  "Server": {  
    "active": 1  
  },  
  "Worker": {  
    "active": 1  
  }  
},  
"startTime": "2019-07-24T08:42:33Z"  
}
```

状态码

表 9-39 状态码

状态码	描述
200	OK
401	Unauthorized
404	Not found
500	Internal error
403	Forbidden

9.1.3.3 查询指定 namespace 下的所有 MXJob

功能介绍

查询Namespace下所有MXJob的详细信息。

URI

GET /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs

表 9-40 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-41 Query 参数

参数	是否必选	描述
fieldSelector	No	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	No	A selector to restrict the list of returned objects by their labels. Defaults to everything.
limit	No	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the continue field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
resourceVersion	No	When specified with a watch call, shows changes that occur after that particular version of a resource. Defaults to changes from the beginning of history. When specified for list: - if unset, then the result is returned from remote storage based on quorum-read flag; - if it's 0, then we simply return what we currently have in cache, no guarantee; - if set to non zero, then the result is at least as fresh as given rv.
timeoutSeconds	No	Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.

参数	是否必选	描述
watch	No	Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-290](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "items": [
    {
      "apiVersion": "kubeflow.org/v1",
      "kind": "MXJob",
      "metadata": {
        "creationTimestamp": "2019-07-24T08:42:33Z",
        "generation": 1,
        "name": "mxnet-job",
        "namespace": "kube-test",
        "resourceVersion": "72476154",
        "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",
        "uid": "fac6dcd2-adee-11e9-8041-340a9837e2a7"
      },
      "spec": {
        "cleanPodPolicy": "Running",
        "jobMode": "MXTrain",
        "mxReplicaSpecs": {
          "Scheduler": {
            "replicas": 1,
            "restartPolicy": "Never",
            "template": {
              "spec": {
                "containers": [
                  {
                    "args": [
                      "-c",
                      "python train_imagenet.py"
                    ],
                    "command": [
                      "/bin/bash"
                    ],
                    "image": "**.215:20202/cci/mxnet:xsw-dis",
                    "name": "mxnet",
                    "resources": {
                      "limits": {
                        "cpu": "1000m",
                        "memory": "2Gi"
                      },
                      "requests": {
                        "cpu": "1000m",
                        "memory": "2Gi"
                      }
                    }
                  }
                ]
              }
            }
          }
        }
      }
    }
  ]
}
```

```
    }
  ],
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ]
}
},
"Server": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "*/.*/.215:20202/cci/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "*/.*/.215:20202/cci/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
```

```
        "cpu": "1000m",
        "memory": "2Gi"
      }
    }
  ],
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ]
}
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T08:42:33Z",
      "lastUpdateTime": "2019-07-24T08:42:33Z",
      "message": "MXJob mxnet-job is created.",
      "reason": "MXJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-24T08:42:33Z",
      "lastUpdateTime": "2019-07-24T08:42:52Z",
      "message": "MXJob mxnet-job is running.",
      "reason": "MXJobRunning",
      "status": "True",
      "type": "Running"
    }
  ],
  "mxReplicaStatuses": {
    "Scheduler": {
      "active": 1
    },
    "Server": {
      "active": 1
    },
    "Worker": {
      "active": 1
    }
  },
  "startTime": "2019-07-24T08:42:33Z"
}
},
"kind": "MXJobList",
"metadata": {
  "continue": "",
  "resourceVersion": "72478902",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs"
}
}
```

状态码

表 9-42 状态码

状态码	描述
200	OK

状态码	描述
401	Unauthorized
404	Not found
500	Internal error

9.1.3.4 删除 namespace 下的所有 MXJob

功能介绍

删除命名空间下的所有MXJob。

URI

DELETE /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs

表 9-43 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-44 Query 参数

参数	是否必选	描述
fieldSelector	No	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	No	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	描述
limit	No	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the continue field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
resourceVersion	No	<p>When specified with a watch call, shows changes that occur after that particular version of a resource. Defaults to changes from the beginning of history. When specified for list: - if unset, then the result is returned from remote storage based on quorum-read flag; - if it's 0, then we simply return what we currently have in cache, no guarantee; - if set to non zero, then the result is at least as fresh as given rv.</p>
timeoutSeconds	No	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>
watch	No	<p>Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.</p>

请求消息

N/A

响应消息

响应参数：

响应参数的详细描述请参见[表6-193](#)。

响应示例：

```
{
  "apiVersion": "kubeflow.org/v1",
  "items": [
    {
      "apiVersion": "kubeflow.org/v1",
      "kind": "MXJob",
      "metadata": {
        "creationTimestamp": "2019-07-24T08:57:01Z",
        "generation": 1,
        "name": "mxnet-job",
        "namespace": "kube-test",
        "resourceVersion": "72481787",
        "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",
        "uid": "001b5d2a-adf1-11e9-ba3a-b44326d0c915"
      },
      "spec": {
        "cleanPodPolicy": "Running",
        "jobMode": "MXTrain",
        "mxReplicaSpecs": {
          "Scheduler": {
            "replicas": 1,
            "restartPolicy": "Never",
            "template": {
              "spec": {
                "containers": [
                  {
                    "args": [
                      "-c",
                      "python train_imagenet.py"
                    ],
                    "command": [
                      "/bin/bash"
                    ],
                    "image": "kubeflow.azurecr.io/mxnet:xsw-dis",
                    "name": "mxnet",
                    "resources": {
                      "limits": {
                        "cpu": "1000m",
                        "memory": "2Gi"
                      },
                      "requests": {
                        "cpu": "1000m",
                        "memory": "2Gi"
                      }
                    }
                  }
                ]
              }
            }
          },
          "imagePullSecrets": [
            {
              "name": "imagepull-secret"
            }
          ]
        }
      }
    }
  ],
  "Server": {
```



```
"replicas": 1,
"restartPolicy": "Never",
"template": {
  "spec": {
    "containers": [
      {
        "args": [
          "-c",
          "python train_imagenet.py"
        ],
        "command": [
          "/bin/bash"
        ],
        "image": "oci://215:20202/cci/mxnet:xsw-dis",
        "name": "mxnet",
        "resources": {
          "limits": {
            "cpu": "1000m",
            "memory": "2Gi"
          },
          "requests": {
            "cpu": "1000m",
            "memory": "2Gi"
          }
        }
      }
    ],
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --
num-epochs 1 --kv-store dist_sync --num-examples 500"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "oci://215:20202/cci/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ],
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ]
    }
  }
}
```

```

    }
  }
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T08:57:01Z",
      "lastUpdateTime": "2019-07-24T08:57:01Z",
      "message": "MXJob mxnet-job is created.",
      "reason": "MXJobCreated",
      "status": "True",
      "type": "Created"
    },
    {
      "lastTransitionTime": "2019-07-24T08:57:01Z",
      "lastUpdateTime": "2019-07-24T08:57:06Z",
      "message": "MXJob mxnet-job is running.",
      "reason": "MXJobRunning",
      "status": "True",
      "type": "Running"
    }
  ],
  "mxReplicaStatuses": {
    "Scheduler": {
      "active": 1
    },
    "Server": {
      "active": 1
    },
    "Worker": {
      "active": 1
    }
  },
  "startTime": "2019-07-24T08:57:01Z"
}
},
"kind": "MXJobList",
"metadata": {
  "continue": "",
  "resourceVersion": "72482111",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs"
}
}

```

状态码

表 3 状态码描述API的状态码。

表 9-45 状态码

状态码	描述
200	OK
401	Unauthorized
500	Internal error

9.1.3.5 删除 MXJob

功能介绍

删除MXJob。

URI

DELETE /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs/{name}

表 9-46 Path 参数

参数	是否必选	描述
name	Yes	name of the MXJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-47 Query 参数

参数	是否必选	描述
dryRun	No	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	No	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.
orphanDependents	No	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the “orphan” finalizer will be added to/removed from the object’s finalizers list. Either this field or PropagationPolicy may be set, but not both.

参数	是否必选	描述
propagationPolicy	No	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	No	If 'true' , then the output is pretty printed.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-193](#)。

响应示例:

```
{
  "kind": "Status",
  "apiVersion": "v1",
  "metadata": {},
  "status": "Success",
  "details": {
    "name": "mxnet-job",
    "group": "kubeflow.org",
    "kind": "mxjobs",
    "uid": "fac6dcd2-adee-11e9-8041-340a9837e2a7"
  }
}
```

状态码

表 9-48 状态码

状态码	描述
200	OK
202	Accepted
401	Unauthorized
500	Internal Error
403	Forbidden

9.1.3.6 更新 MXJob

功能介绍

更新MXJob。如下字段可被更新：

- metadata.labels
- metadata.annotations
- spec.activeDeadlineSeconds
- spec.ttlSecondsAfterFinished
- spec.cleanPodPolicy

URI

PATCH /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs/{name}

表 9-49 Path 参数

参数	是否必选	描述
name	Yes	name of the MXJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-50 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数：

“Content-Type” 消息头说明请参见[PATCH请求方法操作说明](#)。

📖 说明

目前只支持“Merge Patch”。

请求示例：

```
Content-Type: application/merge-patch+json
{
  "metadata": {
    "labels": {
      "app": "test"
    }
  }
}
```

```
}  
}
```

响应消息

响应参数:

响应参数的详细描述请参见[表6-283](#)。

响应示例:

```
{  
  "apiVersion": "kubeflow.org/v1",  
  "kind": "MXJob",  
  "metadata": {  
    "creationTimestamp": "2019-07-24T08:58:47Z",  
    "generation": 1,  
    "labels": {  
      "app": "test"  
    },  
    "name": "mxnet-job",  
    "namespace": "kube-test",  
    "resourceVersion": "72482855",  
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",  
    "uid": "3f304b53-adf1-11e9-8041-340a9837e2a7"  
  },  
  "spec": {  
    "cleanPodPolicy": "Running",  
    "jobMode": "MXTrain",  
    "mxReplicaSpecs": {  
      "Scheduler": {  
        "replicas": 1,  
        "restartPolicy": "Never",  
        "template": {  
          "spec": {  
            "containers": [  
              {  
                "args": [  
                  "-c",  
                  "python train_imagenet.py"  
                ],  
                "command": [  
                  "/bin/bash"  
                ],  
                "image": "*/*.215:20202/cci/mxnet:xsw-dis",  
                "name": "mxnet",  
                "resources": {  
                  "limits": {  
                    "cpu": "1000m",  
                    "memory": "2Gi"  
                  },  
                  "requests": {  
                    "cpu": "1000m",  
                    "memory": "2Gi"  
                  }  
                }  
              }  
            ],  
            "imagePullSecrets": [  
              {  
                "name": "imagepull-secret"  
              }  
            ]  
          }  
        }  
      }  
    },  
    "Server": {  
      "replicas": 1,  
      "restartPolicy": "Never",  
    }  
  }  
}
```

```
"template": {
  "spec": {
    "containers": [
      {
        "args": [
          "-c",
          "python train_imagenet.py"
        ],
        "command": [
          "/bin/bash"
        ],
        "image": "**:*215:20202/ci/mxnet:xsw-dis",
        "name": "mxnet",
        "resources": {
          "limits": {
            "cpu": "1000m",
            "memory": "2Gi"
          },
          "requests": {
            "cpu": "1000m",
            "memory": "2Gi"
          }
        }
      }
    ],
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "**:*215:20202/ci/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ],
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ]
    }
  }
}
```

```
    }
  },
  "status": {
    "conditions": [
      {
        "lastTransitionTime": "2019-07-24T08:58:47Z",
        "lastUpdateTime": "2019-07-24T08:58:47Z",
        "message": "MXJob mxnet-job is created.",
        "reason": "MXJobCreated",
        "status": "True",
        "type": "Created"
      },
      {
        "lastTransitionTime": "2019-07-24T08:58:47Z",
        "lastUpdateTime": "2019-07-24T08:58:52Z",
        "message": "MXJob mxnet-job is running.",
        "reason": "MXJobRunning",
        "status": "True",
        "type": "Running"
      }
    ],
    "mxReplicaStatuses": {
      "Scheduler": {
        "active": 1
      },
      "Server": {
        "active": 1
      },
      "Worker": {
        "active": 1
      }
    },
    "startTime": "2019-07-24T08:58:47Z"
  }
}
```

状态码

表 9-51 状态码

状态码	描述
200	OK
401	Unauthorized
500	Internal Error
403	Forbidden
409	Conflict
400	BadRequest

9.1.3.7 替换 MXJob

功能介绍

替换MXJob。如下字段可被替换：

- metadata.labels
- metadata.annotations
- spec.activeDeadlineSeconds
- spec.ttlSecondsAfterFinished
- spec.cleanPodPolicy

URI

PUT /apis/kubeflow.org/v1/namespaces/{namespace}/mxjobs/{name}

表 9-52 Path 参数

参数	是否必选	描述
name	Yes	name of the MXJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-53 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数:

请求参数的详细描述请参见[表6-283](#)。

请求示例:

更改MXJob的结束存活时间ttlSecondsAfterFinished:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "MXJob",
  "metadata": {
    "creationTimestamp": "2019-07-29T03:34:33Z",
    "generation": 2,
    "name": "mxnet-job",
    "namespace": "kube-test",
    "resourceVersion": "75615427",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",
    "uid": "c82e664f-b1b1-11e9-b310-b44326d0c915"
  },
  "spec": {
    "cleanPodPolicy": "Running",
    "jobMode": "MXTrain",
    "mxReplicaSpecs": {
      "Scheduler": {
        "replicas": 1,
        "restartPolicy": "Never",
        "template": {
```

```
"spec": {
  "containers": [
    {
      "args": [
        "-c",
        "python train_imagenet.py"
      ],
      "command": [
        "/bin/bash"
      ],
      "image": "100.79.1.215:20202/paas_cci_w00427225/mxnet:xsw-dis",
      "name": "mxnet",
      "resources": {
        "limits": {
          "cpu": "1000m",
          "memory": "2Gi"
        },
        "requests": {
          "cpu": "1000m",
          "memory": "2Gi"
        }
      }
    }
  ],
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ]
},
"Server": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "100.79.1.215:20202/paas_cci_w00427225/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ],
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ]
    }
  }
},
"Worker": {
```

```
    "replicas": 1,
    "restartPolicy": "Never",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "-c",
              "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
            ],
            "command": [
              "/bin/bash"
            ],
            "image": "100.79.1.215:20202/paas_cci_w00427225/mxnet:xsw-dis",
            "name": "mxnet",
            "resources": {
              "limits": {
                "cpu": "1000m",
                "memory": "2Gi"
              },
              "requests": {
                "cpu": "1000m",
                "memory": "2Gi"
              }
            }
          }
        ],
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      }
    },
    "ttlSecondsAfterFinished": 10000
  },
  "status": {
    "conditions": [
      {
        "lastTransitionTime": "2019-07-29T03:34:33Z",
        "lastUpdateTime": "2019-07-29T03:34:33Z",
        "message": "MXJob mxnet-job is created.",
        "reason": "MXJobCreated",
        "status": "True",
        "type": "Created"
      },
      {
        "lastTransitionTime": "2019-07-29T03:34:33Z",
        "lastUpdateTime": "2019-07-29T03:35:22Z",
        "message": "MXJob mxnet-job is running.",
        "reason": "MXJobRunning",
        "status": "True",
        "type": "Running"
      }
    ],
    "mxReplicaStatuses": {
      "Scheduler": {
        "active": 1
      },
      "Server": {
        "active": 1
      },
      "Worker": {
        "active": 1
      }
    }
  }
},
```

```
    "startTime": "2019-07-29T03:34:33Z"  
  }  
}
```

响应消息

响应参数:

响应参数的详细描述请参见[表6-283](#)。

响应示例:

```
{  
  "apiVersion": "kubeflow.org/v1",  
  "kind": "MXJob",  
  "metadata": {  
    "creationTimestamp": "2019-07-29T03:34:33Z",  
    "generation": 2,  
    "name": "mxnet-job",  
    "namespace": "kube-test",  
    "resourceVersion": "75615427",  
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/mxjobs/mxnet-job",  
    "uid": "c82e664f-b1b1-11e9-b310-b44326d0c915"  
  },  
  "spec": {  
    "cleanPodPolicy": "Running",  
    "jobMode": "MXTrain",  
    "mxReplicaSpecs": {  
      "Scheduler": {  
        "replicas": 1,  
        "restartPolicy": "Never",  
        "template": {  
          "spec": {  
            "containers": [  
              {  
                "args": [  
                  "-c",  
                  "python train_imagenet.py"  
                ],  
                "command": [  
                  "/bin/bash"  
                ],  
                "image": "100.79.1.215:20202/paas_cci_w00427225/mxnet:xsw-dis",  
                "name": "mxnet",  
                "resources": {  
                  "limits": {  
                    "cpu": "1000m",  
                    "memory": "2Gi"  
                  },  
                  "requests": {  
                    "cpu": "1000m",  
                    "memory": "2Gi"  
                  }  
                }  
              }  
            ],  
            "imagePullSecrets": [  
              {  
                "name": "imagepull-secret"  
              }  
            ]  
          }  
        }  
      }  
    }  
  },  
  "Server": {  
    "replicas": 1,  
    "restartPolicy": "Never",  
    "template": {  
      "spec": {
```

```
"containers": [
  {
    "args": [
      "-c",
      "python train_imagenet.py"
    ],
    "command": [
      "/bin/bash"
    ],
    "image": "100.79.1.215:20202/paas_cci_w00427225/mxnet:xsw-dis",
    "name": "mxnet",
    "resources": {
      "limits": {
        "cpu": "1000m",
        "memory": "2Gi"
      },
      "requests": {
        "cpu": "1000m",
        "memory": "2Gi"
      }
    }
  }
],
"imagePullSecrets": [
  {
    "name": "imagepull-secret"
  }
]
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "Never",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "-c",
            "python train_imagenet.py --benchmark 1 --network resnet --batch-size 1 --num-epochs 1 --kv-store dist_sync --num-examples 500"
          ],
          "command": [
            "/bin/bash"
          ],
          "image": "100.79.1.215:20202/paas_cci_w00427225/mxnet:xsw-dis",
          "name": "mxnet",
          "resources": {
            "limits": {
              "cpu": "1000m",
              "memory": "2Gi"
            },
            "requests": {
              "cpu": "1000m",
              "memory": "2Gi"
            }
          }
        }
      ],
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ]
    }
  }
},
}
```

```

        "ttlSecondsAfterFinished": 10000
    },
    "status": {
      "conditions": [
        {
          "lastTransitionTime": "2019-07-29T03:34:33Z",
          "lastUpdateTime": "2019-07-29T03:34:33Z",
          "message": "MXJob mxnet-job is created.",
          "reason": "MXJobCreated",
          "status": "True",
          "type": "Created"
        },
        {
          "lastTransitionTime": "2019-07-29T03:34:33Z",
          "lastUpdateTime": "2019-07-29T03:35:22Z",
          "message": "MXJob mxnet-job is running.",
          "reason": "MXJobRunning",
          "status": "True",
          "type": "Running"
        }
      ]
    },
    "mxReplicaStatuses": {
      "Scheduler": {
        "active": 1
      },
      "Server": {
        "active": 1
      },
      "Worker": {
        "active": 1
      }
    },
    "startTime": "2019-07-29T03:34:33Z"
  }
}

```

状态码

表 9-54 状态码

状态码	描述
200	OK
201	Created
401	Unauthorized
400	BadRequest
500	Internal Error
403	Forbidden

9.1.4 PyTorchJob

9.1.4.1 创建 PyTorchJob

功能介绍

创建PyTorchJob。

PyTorchJob即PyTorch任务，是基于PyTorch开源框架的kubernetes自定义资源类型，有多种角色可以配置，能够帮助我们更简单地实现PyTorch的训练。

URI

POST /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs

表 9-55 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-56 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数：

请求参数的详细描述请参见[表6-286](#)。

请求示例：

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "PyTorchJob",
  "metadata": {
    "name": "pytorch-test"
  },
  "spec": {
    "pytorchReplicaSpecs": {
      "Master": {
        "replicas": 1,
        "restartPolicy": "Never",
        "template": {
          "spec": {
            "containers": [
              {
                "name": "pytorch",
                "image": "**:*215:20202/gcs/pytorch-cpu:v1",
                "command": [
                  "python",
                  "/var/mnist.py"
                ],
                "args": [
                  "--backend",

```

```
        "gloo"
      ],
      "resources": {
        "limits": {
          "cpu": 2,
          "memory": "4Gi"
        },
        "requests": {
          "cpu": 2,
          "memory": "4Gi"
        }
      }
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "OnFailure",
  "template": {
    "spec": {
      "containers": [
        {
          "name": "pytorch",
          "image": "***.215:20202/gcs/pytorch-cpu:v1",
          "command": [
            "python",
            "/var/mnist.py"
          ],
          "args": [
            "--backend",
            "gloo"
          ],
          "resources": {
            "limits": {
              "cpu": 2,
              "memory": "4Gi"
            },
            "requests": {
              "cpu": 2,
              "memory": "4Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
}
}
```

响应消息

响应参数:

响应参数的详细描述请参见[表6-286](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "PyTorchJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T10:29:45Z",
    "generation": 1,
    "name": "pytorch-test",
    "namespace": "kube-test",
    "resourceVersion": "72516798",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
    "uid": "f4c79668-adfd-11e9-8041-340a9837e2a7"
  },
  "spec": {
    "pytorchReplicaSpecs": {
      "Master": {
        "replicas": 1,
        "restartPolicy": "Never",
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "--backend",
                  "gloo"
                ],
                "command": [
                  "python",
                  "/var/mnist.py"
                ],
                "image": "**.*.215:20202/gcs/pytorch-cpu:v1",
                "name": "pytorch",
                "resources": {
                  "limits": {
                    "cpu": 2,
                    "memory": "4Gi"
                  },
                  "requests": {
                    "cpu": 2,
                    "memory": "4Gi"
                  }
                }
              }
            ]
          }
        },
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      }
    },
    "Worker": {
      "replicas": 1,
      "restartPolicy": "OnFailure",
      "template": {
        "spec": {
          "containers": [
            {
              "args": [
                "--backend",
                "gloo"
              ],
              "command": [
                "python",
                "/var/mnist.py"
              ],
              "image": "**.*.215:20202/gcs/pytorch-cpu:v1",
              "name": "pytorch",
            }
          ]
        }
      }
    }
  }
}
```

```
        "resources": {
          "limits": {
            "cpu": 2,
            "memory": "4Gi"
          },
          "requests": {
            "cpu": 2,
            "memory": "4Gi"
          }
        }
      ],
      "imagePullSecrets": [
        {
          "name": "imagepull-secret"
        }
      ]
    }
  },
  "status": {
  }
}
```

状态码

表 9-57 状态码

状态码	描述
200	OK
201	Created
202	Accepted
401	Unauthorized
400	Badrequest
500	Internal error
403	Forbidden

9.1.4.2 查询 PyTorchJob

功能介绍

查询PyTorchJob的详细信息。

URI

GET /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs/{name}

表 9-58 Path 参数

参数	是否必选	描述
name	Yes	name of the PyTorchJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-59 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-286](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "PyTorchJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T10:29:45Z",
    "generation": 1,
    "name": "pytorch-test",
    "namespace": "kube-test",
    "resourceVersion": "72516798",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
    "uid": "f4c79668-adfd-11e9-8041-340a9837e2a7"
  },
  "spec": {
    "pytorchReplicaSpecs": {
      "Master": {
        "replicas": 1,
        "restartPolicy": "Never",
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "--backend",
                  "gloo"
                ],
                "command": [
                  "python",
                  "/var/mnist.py"
                ],
                "image": "**.215:20202/gcs/pytorch-cpu:v1",
                "name": "pytorch",
                "resources": {
                  "limits": {
```

```
        "cpu": 2,
        "memory": "4Gi"
      },
      "requests": {
        "cpu": 2,
        "memory": "4Gi"
      }
    }
  ],
  "imagePullSecrets": [
    {
      "name": "imagepull-secret"
    }
  ]
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "OnFailure",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "--backend",
            "gloo"
          ],
          "command": [
            "python",
            "/var/mnist.py"
          ],
          "image": "**:*215:20202/gcs/pytorch-cpu:v1",
          "name": "pytorch",
          "resources": {
            "limits": {
              "cpu": 2,
              "memory": "4Gi"
            },
            "requests": {
              "cpu": 2,
              "memory": "4Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T10:30:34Z",
      "lastUpdateTime": "2019-07-24T10:30:34Z",
      "message": "PyTorchJob pytorch-test is created.",
      "reason": "PyTorchJobCreated",
      "status": "True",
      "type": "Created"
    }
  ],
  "replicaStatuses": {
```

```
    "Master": {}  
  },  
  "startTime": "2019-07-24T10:30:34Z"  
}
```

状态码

表 9-60 状态码

状态码	描述
200	OK
401	Unauthorized
404	Not found
500	Internal error
403	Forbidden

9.1.4.3 查询指定 namespace 下的所有 PyTorchJob

功能介绍

查询Namespace下所有PyTorchJob的详细信息。

URI

GET /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs

表 9-61 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-62 Query 参数

参数	是否必选	描述
fieldSelector	No	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	No	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	描述
limit	No	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the continue field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
resourceVersion	No	<p>When specified with a watch call, shows changes that occur after that particular version of a resource. Defaults to changes from the beginning of history. When specified for list: - if unset, then the result is returned from remote storage based on quorum-read flag; - if it's 0, then we simply return what we currently have in cache, no guarantee; - if set to non zero, then the result is at least as fresh as given rv.</p>
timeoutSeconds	No	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>
watch	No	<p>Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.</p>

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-291](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "items": [
    {
      "apiVersion": "kubeflow.org/v1",
      "kind": "PyTorchJob",
      "metadata": {
        "creationTimestamp": "2019-07-24T10:29:45Z",
        "generation": 1,
        "name": "pytorch-test",
        "namespace": "kube-test",
        "resourceVersion": "72516798",
        "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
        "uid": "f4c79668-adfd-11e9-8041-340a9837e2a7"
      },
      "spec": {
        "pytorchReplicaSpecs": {
          "Master": {
            "replicas": 1,
            "restartPolicy": "Never",
            "template": {
              "spec": {
                "containers": [
                  {
                    "args": [
                      "--backend",
                      "gloo"
                    ],
                    "command": [
                      "python",
                      "/var/mnist.py"
                    ],
                    "image": "*/*.215:20202/gcs/pytorch-cpu:v1",
                    "name": "pytorch",
                    "resources": {
                      "limits": {
                        "cpu": 2,
                        "memory": "4Gi"
                      },
                      "requests": {
                        "cpu": 2,
                        "memory": "4Gi"
                      }
                    }
                  }
                ]
              }
            }
          },
          "imagePullSecrets": [
            {
              "name": "imagepull-secret"
            }
          ]
        }
      },
      "Worker": {
        "replicas": 1,
```

```
    "restartPolicy": "OnFailure",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "--backend",
              "gloo"
            ],
            "command": [
              "python",
              "/var/mnist.py"
            ],
            "image": "gcr.io/ml-org:20202/gcs/pytorch-cpu:v1",
            "name": "pytorch",
            "resources": {
              "limits": {
                "cpu": 2,
                "memory": "4Gi"
              },
              "requests": {
                "cpu": 2,
                "memory": "4Gi"
              }
            }
          }
        ],
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      }
    }
  },
  "status": {
    "conditions": [
      {
        "lastTransitionTime": "2019-07-24T10:30:34Z",
        "lastUpdateTime": "2019-07-24T10:30:34Z",
        "message": "PyTorchJob pytorch-test is created.",
        "reason": "PyTorchJobCreated",
        "status": "True",
        "type": "Created"
      }
    ],
    "replicaStatuses": {
      "Master": {}
    },
    "startTime": "2019-07-24T10:30:34Z"
  }
},
"kind": "PyTorchJobList",
"metadata": {
  "continue": "",
  "resourceVersion": "72517974",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs"
}
}
```


状态码

表 9-63 状态码

状态码	描述
200	OK
401	Unauthorized
404	Not found
500	Internal error

9.1.4.4 删除 namespace 下的所有 PyTorchJob

功能介绍

删除命名空间下的所有PyTorchJob。

URI

DELETE /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs

表 9-64 Path 参数

参数	是否必选	描述
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-65 Query 参数

参数	是否必选	描述
fieldSelector	No	A selector to restrict the list of returned objects by their fields. Defaults to everything.
labelSelector	No	A selector to restrict the list of returned objects by their labels. Defaults to everything.

参数	是否必选	描述
limit	No	<p>limit is a maximum number of responses to return for a list call. If more items exist, the server will set the continue field on the list metadata to a value that can be used with the same initial query to retrieve the next set of results. Setting a limit may return fewer than the requested amount of items (up to zero items) in the event all requested objects are filtered out and clients should only use the presence of the continue field to determine whether more results are available. Servers may choose not to support the limit argument and will return all of the available results. If limit is specified and the continue field is empty, clients may assume that no more results are available. This field is not supported if watch is true.</p> <p>The server guarantees that the objects returned when using continue will be identical to issuing a single list call without a limit - that is, no objects created, modified, or deleted after the first request is issued will be included in any subsequent continued requests. This is sometimes referred to as a consistent snapshot, and ensures that a client that is using limit to receive smaller chunks of a very large result can ensure they see all possible objects. If objects are updated during a chunked list the version of the object that was present at the time the first list result was calculated is returned.</p>
resourceVersion	No	<p>When specified with a watch call, shows changes that occur after that particular version of a resource. Defaults to changes from the beginning of history. When specified for list: - if unset, then the result is returned from remote storage based on quorum-read flag; - if it's 0, then we simply return what we currently have in cache, no guarantee; - if set to non zero, then the result is at least as fresh as given rv.</p>
timeoutSeconds	No	<p>Timeout for the list/watch call. This limits the duration of the call, regardless of any activity or inactivity.</p>
watch	No	<p>Watch for changes to the described resources and return them as a stream of add, update, and remove notifications. Specify resourceVersion.</p>

请求消息

N/A

响应消息

响应参数：

响应参数的详细描述请参见[表6-291](#)。

响应示例：

```
{
  "apiVersion": "kubeflow.org/v1",
  "items": [
    {
      "apiVersion": "kubeflow.org/v1",
      "kind": "PyTorchJob",
      "metadata": {
        "creationTimestamp": "2019-07-24T10:29:45Z",
        "generation": 1,
        "name": "pytorch-test",
        "namespace": "kube-test",
        "resourceVersion": "72516798",
        "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
        "uid": "f4c79668-adfd-11e9-8041-340a9837e2a7"
      },
      "spec": {
        "pytorchReplicaSpecs": {
          "Master": {
            "replicas": 1,
            "restartPolicy": "Never",
            "template": {
              "spec": {
                "containers": [
                  {
                    "args": [
                      "--backend",
                      "gloo"
                    ],
                    "command": [
                      "python",
                      "/var/mnist.py"
                    ],
                    "image": "*/*.215:20202/gcs/pytorch-cpu:v1",
                    "name": "pytorch",
                    "resources": {
                      "limits": {
                        "cpu": 2,
                        "memory": "4Gi"
                      },
                      "requests": {
                        "cpu": 2,
                        "memory": "4Gi"
                      }
                    }
                  }
                ]
              }
            }
          },
          "imagePullSecrets": [
            {
              "name": "imagepull-secret"
            }
          ]
        }
      }
    },
    "Worker": {
      "replicas": 1,
```

```
    "restartPolicy": "OnFailure",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "--backend",
              "gloo"
            ],
            "command": [
              "python",
              "/var/mnist.py"
            ],
            "image": "**:*215:20202/gcs/pytorch-cpu:v1",
            "name": "pytorch",
            "resources": {
              "limits": {
                "cpu": 2,
                "memory": "4Gi"
              },
              "requests": {
                "cpu": 2,
                "memory": "4Gi"
              }
            }
          }
        ],
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      }
    }
  },
  "status": {
    "conditions": [
      {
        "lastTransitionTime": "2019-07-24T10:30:34Z",
        "lastUpdateTime": "2019-07-24T10:30:34Z",
        "message": "PyTorchJob pytorch-test is created.",
        "reason": "PyTorchJobCreated",
        "status": "True",
        "type": "Created"
      }
    ],
    "replicaStatuses": {
      "Master": {}
    },
    "startTime": "2019-07-24T10:30:34Z"
  }
},
"kind": "PyTorchJobList",
"metadata": {
  "continue": "",
  "resourceVersion": "72518429",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs"
}
}
```

状态码

表 9-66 状态码

状态码	描述
200	OK
401	Unauthorized
500	Internal error

9.1.4.5 删除 PyTorchJob

功能介绍

删除PyTorchJob。

URI

DELETE /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs/{name}

表 9-67 Path 参数

参数	是否必选	描述
name	Yes	name of the PyTorchJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-68 Query 参数

参数	是否必选	描述
dryRun	No	When present, indicates that modifications should not be persisted. An invalid or unrecognized dryRun directive will result in an error response and no further processing of the request. Valid values are: - All: all dry run stages will be processed
gracePeriodSeconds	No	The duration in seconds before the object should be deleted. Value must be non-negative integer. The value zero indicates delete immediately. If this value is nil, the default grace period for the specified type will be used. Defaults to a per object value if not specified. zero means delete immediately.

参数	是否必选	描述
orphanDependents	No	Deprecated: please use the PropagationPolicy, this field will be deprecated in 1.7. Should the dependent objects be orphaned. If true/false, the "orphan" finalizer will be added to/removed from the object's finalizers list. Either this field or PropagationPolicy may be set, but not both.
propagationPolicy	No	Whether and how garbage collection will be performed. Either this field or OrphanDependents may be set, but not both. The default policy is decided by the existing finalizer set in the metadata.finalizers and the resource-specific default policy. Acceptable values are: 'Orphan' - orphan the dependents; 'Background' - allow the garbage collector to delete the dependents in the background; 'Foreground' - a cascading policy that deletes all dependents in the foreground.
pretty	No	If 'true', then the output is pretty printed.

请求消息

N/A

响应消息

响应参数:

响应参数的详细描述请参见[表6-193](#)。

响应示例:

```
{
  "kind": "Status",
  "apiVersion": "v1",
  "metadata": {},
  "status": "Success",
  "details": {
    "name": "pytorch-test",
    "group": "kubeflow.org",
    "kind": "pytorchjobs",
    "uid": "be7696f9-adfe-11e9-aaa4-340a9837e413"
  }
}
```

状态码

表 9-69 状态码

状态码	描述
200	OK

状态码	描述
202	Accepted
401	Unauthorized
500	Internal Error
403	Forbidden

9.1.4.6 更新 PyTorchJob

功能介绍

更新PyTorchJob。如下字段可被更新：

- metadata.labels
- metadata.annotations
- spec.activeDeadlineSeconds
- spec.ttlSecondsAfterFinished
- spec.cleanPodPolicy

URI

PATCH /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs/{name}

表 9-70 Path 参数

参数	是否必选	描述
name	Yes	name of the PyTorchJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-71 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数：

“Content-Type” 消息头说明请参见[PATCH请求方法操作说明](#)。

📖 说明

目前只支持“Merge Patch”。

请求示例:

```
Content-Type: application/merge-patch+json
{
  "metadata": {
    "labels": {
      "app": "test"
    }
  }
}
```

响应消息

响应参数:

响应参数的详细描述请参见[表6-286](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "PyTorchJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T10:35:38Z",
    "generation": 1,
    "labels": {
      "app": "test"
    }
  },
  "name": "pytorch-test",
  "namespace": "kube-test",
  "resourceVersion": "72519229",
  "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
  "uid": "c6e548f1-adfe-11e9-ba3a-b44326d0c915"
},
"spec": {
  "pytorchReplicaSpecs": {
    "Master": {
      "replicas": 1,
      "restartPolicy": "Never",
      "template": {
        "spec": {
          "containers": [
            {
              "args": [
                "--backend",
                "gloo"
              ],
              "command": [
                "python",
                "/var/mnist.py"
              ],
              "image": "***.215:20202/gcs/pytorch-cpu:v1",
              "name": "pytorch",
              "resources": {
                "limits": {
                  "cpu": 2,
                  "memory": "4Gi"
                },
                "requests": {
                  "cpu": 2,
                  "memory": "4Gi"
                }
              }
            }
          ]
        }
      }
    }
  }
}
```



```
    ],
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
},
"Worker": {
  "replicas": 1,
  "restartPolicy": "OnFailure",
  "template": {
    "spec": {
      "containers": [
        {
          "args": [
            "--backend",
            "gloo"
          ],
          "command": [
            "python",
            "/var/mnist.py"
          ],
          "image": "**.*.215:20202/gcs/pytorch-cpu:v1",
          "name": "pytorch",
          "resources": {
            "limits": {
              "cpu": 2,
              "memory": "4Gi"
            },
            "requests": {
              "cpu": 2,
              "memory": "4Gi"
            }
          }
        }
      ]
    },
    "imagePullSecrets": [
      {
        "name": "imagepull-secret"
      }
    ]
  }
}
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T10:36:26Z",
      "lastUpdateTime": "2019-07-24T10:36:26Z",
      "message": "PyTorchJob pytorch-test is created.",
      "reason": "PyTorchJobCreated",
      "status": "True",
      "type": "Created"
    }
  ],
  "replicaStatuses": {
    "Master": {}
  },
  "startTime": "2019-07-24T10:36:26Z"
}
```

状态码

表 9-72 状态码

状态码	描述
200	OK
401	Unauthorized
500	Internal Error
403	Forbidden
409	Conflict
400	BadRequest

9.1.4.7 替换 PyTorchJob

功能介绍

替换PyTorchJob。如下字段可被替换：

- metadata.labels
- metadata.annotations
- spec.activeDeadlineSeconds
- spec.ttlSecondsAfterFinished
- spec.cleanPodPolicy

URI

PUT /apis/kubeflow.org/v1/namespaces/{namespace}/pytorchjobs/{name}

表 9-73 Path 参数

参数	是否必选	描述
name	Yes	name of the PyTorchJob
namespace	Yes	object name and auth scope, such as for teams and projects

表 9-74 Query 参数

参数	是否必选	描述
pretty	No	If 'true' , then the output is pretty printed.

请求消息

请求参数:

请求参数的详细描述请参见[表6-286](#)。

请求示例:

更改PyTorchJob的结束存活时间ttlSecondsAfterFinished:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "PyTorchJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T10:35:38Z",
    "generation": 2,
    "labels": {
      "app": "test"
    },
    "name": "pytorch-test",
    "namespace": "kube-test",
    "resourceVersion": "72519846",
    "selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
    "uid": "c6e548f1-adfe-11e9-ba3a-b44326d0c915"
  },
  "spec": {
    "pytorchReplicaSpecs": {
      "Master": {
        "replicas": 1,
        "restartPolicy": "Never",
        "template": {
          "spec": {
            "containers": [
              {
                "args": [
                  "--backend",
                  "gloo"
                ],
                "command": [
                  "python",
                  "/var/mnist.py"
                ],
                "image": "**:*215:20202/gcs/pytorch-cpu:v1",
                "name": "pytorch",
                "resources": {
                  "limits": {
                    "cpu": 2,
                    "memory": "4Gi"
                  },
                  "requests": {
                    "cpu": 2,
                    "memory": "4Gi"
                  }
                }
              }
            ]
          }
        },
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      },
      "Worker": {
        "replicas": 1,
        "restartPolicy": "OnFailure",
        "template": {
          "spec": {
```

```
"containers": [
  {
    "args": [
      "--backend",
      "gloo"
    ],
    "command": [
      "python",
      "/var/mnist.py"
    ],
    "image": "**.*.215:20202/gcs/pytorch-cpu:v1",
    "name": "pytorch",
    "resources": {
      "limits": {
        "cpu": 2,
        "memory": "4Gi"
      },
      "requests": {
        "cpu": 2,
        "memory": "4Gi"
      }
    }
  }
],
"imagePullSecrets": [
  {
    "name": "imagepull-secret"
  }
]
}
},
"ttlSecondsAfterFinished": 3000
},
"status": {
  "conditions": [
    {
      "lastTransitionTime": "2019-07-24T10:36:26Z",
      "lastUpdateTime": "2019-07-24T10:36:26Z",
      "message": "PyTorchJob pytorch-test is created.",
      "reason": "PyTorchJobCreated",
      "status": "True",
      "type": "Created"
    }
  ],
  "replicaStatuses": {
    "Master": {}
  },
  "startTime": "2019-07-24T10:36:26Z"
}
}
```

响应消息

响应参数:

响应参数的详细描述请参见[表6-286](#)。

响应示例:

```
{
  "apiVersion": "kubeflow.org/v1",
  "kind": "PyTorchJob",
  "metadata": {
    "creationTimestamp": "2019-07-24T10:35:38Z",
    "generation": 2,
    "labels": {
      "app": "test"
    }
  }
}
```

```
},
"name": "pytorch-test",
"namespace": "kube-test",
"resourceVersion": "72519846",
"selfLink": "/apis/kubeflow.org/v1/namespaces/kube-test/pytorchjobs/pytorch-test",
"uid": "c6e548f1-adfe-11e9-ba3a-b44326d0c915"
},
"spec": {
  "pytorchReplicaSpecs": {
    "Master": {
      "replicas": 1,
      "restartPolicy": "Never",
      "template": {
        "spec": {
          "containers": [
            {
              "args": [
                "--backend",
                "gloo"
              ],
              "command": [
                "python",
                "/var/mnist.py"
              ],
              "image": "gcr.io/ml-org/pytorch-cpu:v1",
              "name": "pytorch",
              "resources": {
                "limits": {
                  "cpu": 2,
                  "memory": "4Gi"
                },
                "requests": {
                  "cpu": 2,
                  "memory": "4Gi"
                }
              }
            }
          ]
        },
        "imagePullSecrets": [
          {
            "name": "imagepull-secret"
          }
        ]
      }
    }
  },
  "Worker": {
    "replicas": 1,
    "restartPolicy": "OnFailure",
    "template": {
      "spec": {
        "containers": [
          {
            "args": [
              "--backend",
              "gloo"
            ],
            "command": [
              "python",
              "/var/mnist.py"
            ],
            "image": "gcr.io/ml-org/pytorch-cpu:v1",
            "name": "pytorch",
            "resources": {
              "limits": {
                "cpu": 2,
                "memory": "4Gi"
              },
              "requests": {
```

```
        "cpu": 2,  
        "memory": "4Gi"  
      }  
    },  
    ],  
    "imagePullSecrets": [  
      {  
        "name": "imagepull-secret"  
      }  
    ]  
  }  
},  
"ttlSecondsAfterFinished": 3000  
},  
"status": {  
  "conditions": [  
    {  
      "lastTransitionTime": "2019-07-24T10:36:26Z",  
      "lastUpdateTime": "2019-07-24T10:36:26Z",  
      "message": "PyTorchJob pytorch-test is created.",  
      "reason": "PyTorchJobCreated",  
      "status": "True",  
      "type": "Created"  
    }  
  ],  
  "replicaStatuses": {  
    "Master": {}  
  },  
  "startTime": "2019-07-24T10:36:26Z"  
}
```

状态码

表 9-75 状态码

状态码	描述
200	OK
201	Created
401	Unauthorized
400	BadRequest
500	Internal Error
403	Forbidden

A 修订记录

发布日期	修订记录
2022-06-30	第一次正式发布